"Sometimes Less is More": Multi-Perspective Exploration of Disclosure Abstractions in Location-Aware Social Mobile Applications

Karen P. Tang

December 2010
CMU-HCII-10-108

Human-Computer Interaction Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:
Jason I. Hong (co-chair), Human-Computer Interaction Institute
Daniel P. Siewiorek (co-chair), Human-Computer Interaction Institute
Anind K. Dey, Human-Computer Interaction Institute
Laura A. Dabbish, Heinz College and Human-Computer Interaction Institute
Lorrie Faith Cranor, Institute for Software Research and Engineering & Public Policy

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

1. REPORT DATE **DEC 2010**	2. REPORT TYPE	3. DATES COVERED **00-00-2010 to 00-00-2010**
4. TITLE AND SUBTITLE **'Sometimes Less is More': Multi-Perspective Exploration of Disclosure Abstractions in Location-Aware Social Mobile Applications**		5a. CONTRACT NUMBER
		5b. GRANT NUMBER
		5c. PROGRAM ELEMENT NUMBER
6. AUTHOR(S)		5d. PROJECT NUMBER
		5e. TASK NUMBER
		5f. WORK UNIT NUMBER
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) **Carnegie Mellon University,Human-Computer Interaction Institute,School of Computer Science,Pittsburgh,PA,15213**		8. PERFORMING ORGANIZATION REPORT NUMBER
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION/AVAILABILITY STATEMENT **Approved for public release; distribution unlimited**

13. SUPPLEMENTARY NOTES

14. ABSTRACT

In the past few years, there has been increasing interest in deploying social location-sharing applications (LSAs) that enable users to continuously sense, collect, and share their location information with others. Yet, despite all the attention LSAs are receiving, studies have found that only a small percentage of mobile consumers actively use these services. One often-cited adoption barrier is that many LSAs do not adequately address end-user privacy concerns for sharing location data. One way to address these privacy concerns is to incorporate support for disclosure abstractions in LSAs. These abstractions provide a middle-ground compromise that provides some degree of privacy protection for end-users, as well as some degree of social value to the users who are consuming the location information. In this dissertation, we look at two specific kinds of abstractions: geographic abstractions (which provide spatial blurring of one?s location) and semantic abstractions (which provide obfuscation by referring to the type of location a place is, rather than by its geographical coordinates). We present results from several studies that examine these abstractions at four different stages: how users reason about location sharing, how users configure their privacy preferences, how users interpret visual representations of their location, and what kinds of outcomes can be expected from users that share abstractions. Based on these studies, we provide empirical evidence that relatively simple privacy mechanisms like disclosure abstractions can simplify rule-based privacy configurations and increase the likelihood of location sharing, though there is still a significant chance that abstractions can be reverse-engineered. Based on qualitative user feedback, we also present several privacy implications for visualizing location information as well. By studying these issues with different types of location sharing applications as well as different user study methodologies, we provide a multi-perspective exploration of end-user privacy concerns regarding general location sharing behaviors for context-aware social mobile applications.

15. SUBJECT TERMS

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT **unclassified**	b. ABSTRACT **unclassified**	c. THIS PAGE **unclassified**	**Same as Report (SAR)**	**176**	

Abstract

In the past few years, there has been increasing interest in deploying social location-sharing applications (LSAs) that enable users to continuously sense, collect, and share their location information with others. Yet, despite all the attention LSAs are receiving, studies have found that only a small percentage of mobile consumers actively use these services. One often-cited adoption barrier is that many LSAs do not adequately address end-user privacy concerns for sharing location data.

One way to address these privacy concerns is to incorporate support for disclosure abstractions in LSAs. These abstractions provide a middle-ground compromise that provides some degree of privacy protection for end-users, as well as some degree of social value to the users who are consuming the location information. In this dissertation, we look at two specific kinds of abstractions: geographic abstractions (which provide spatial blurring of one's location) and semantic abstractions (which provide obfuscation by referring to the type of location a place is, rather than by its geographical coordinates).

We present results from several studies that examine these abstractions at four different stages: how users reason about location sharing, how users configure their privacy preferences, how users interpret visual representations of their location, and what kinds of outcomes can be expected from users that share abstractions. Based on these studies, we provide empirical evidence that relatively simple privacy mechanisms like disclosure abstractions can simplify rule-based privacy configurations and increase the likelihood of location sharing, though there is still a significant chance that abstractions can be reverse-engineered. Based on qualitative user feedback, we also present several privacy implications for visualizing location information as well. By studying these issues with

different types of location sharing applications as well as different user study methodologies, we provide a multi-perspective exploration of end-user privacy concerns regarding general location sharing behaviors for context-aware social mobile applications.

Acknowledgements

First and foremost, I must thank my advisors Jason and Dan. They have been extremely patient with me throughout my time in graduate school. I have dabbled in a series of research topics before I settled on this dissertation topic and they have graciously allowed me a lot of freedom to explore my own ideas in this space. They have been wonderful mentors and have taught me many things about research and what is means to be an effective researcher. The journey has not always been easy or smooth, but they have always been there to encourage and guide me through the obstacles.

I also thank my other committee members, Anind Dey, Laura Dabbish, and Lorrie Cranor, for their comments and suggestions that have undoubtedly improved this dissertation. It took me longer than I thought to craft the proposal and it still needed several revisions to reach its current state. Without their help and constructive criticism, this dissertation would not what it is today. They have also all provided me with great advice about career and academia and I am appreciative that I can get such honest, yet encouraging, guidance from them.

To my parents, I owe my interest in computer science. While the decision to go to graduate school was initially a surprise to them, they have since remained supportive of my pursuit of graduate studies and have been incredibly patient during the long journey.

 I also owe a big thanks to the HCII faculty who have provided me with a stimulating and challenging research environment to which I am proud to be part of. They have shown the importance and success of interdisciplinary research and I am hopeful that I can carry on this research spirit in my future endeavors. I would also like to thank the HCII support

to my parents,
for their lifelong support and patience

Table of Contents

List of Figures

List of Tables

1. Introduction: The Era of Location Apps

Due to the recent E911 mandate [Federal Communications Commission, 2000] and advancements in mobile hardware sensing, It is now commonplace to find highly-accurate positioning technology, like global positioning systems (GPS), embedded in today's mobile phones [Meyer, 2008; Zahradnik, 2009]. The ubiquity of such location-aware devices has led to an abundance of location-based services (LBSs), much of which has been driven by a growing interest from industry (see Figure 1). In June 2009, a total of roughly 45,000 iPhone applications were available for download from Apple's iPhone App Store [Apple, 2008] and 2,800 of these were location-based [Skyhook Wireless, 2009a]. Since then, the total number of iPhone applications has continued to increase, reaching over 250,000 [148apps, 2010]. Of particular interest is that location-based services have maintained a similar increasing trend as well; as of February 2010, there were over 6,000 location-based applications [Skyhook Wireless, 2003], demonstrating a more than two-fold increase in an eight-month time period. There are also several location-based services being released on Google's Android Market (see Figure 1, right). As of February 2010, there were over 1,000 location-based applications available, representing about 5% of the total mobile applications on the Android Market [Skyhook Wireless, 2003].

The push for location-based applications signifies a continuing realization towards ubiquitous computing [Weiser, 1991] and offers several benefits for end-users including coordination (e,g., [Colbert, 2001]), navigation & wayfinding (e.g., [May, Ross, Bayer, and Tarkiainen, 2003]), and location-based local searches (e.g., [Mokbel and Aref, 2006; Sohn, Li, Griswold, and Hollan, 2008]). However, in recent years, a new class of applications has emerged, where location-based services are merging with online social

Figure 1. (left) Number of location-based services (LBSs) made available through Apple's iPhone App Store, between June 2008 and February 2010 [Skyhook Wireless, 2003]. The increasing trend of LBSs suggests a growing interest towards developing location-aware applications. (right) Number of location-based services made available through Google's Android Market, between October 2008 and February 2010 [Skyhook Wireless, 2003]. We see similar increasing trends for both the Android phone and Apple's iPhone.

networks. In fact, as of May 2009, social networking is the second most popular type of location-based service being developed for mobile phones [Skyhook Wireless, 2009b]. By leveraging a user's social networks, location-based services are moving towards supporting *social location sharing*, as opposed to simply using location information to support service transactions like search-related tasks [Wikipedia, 2001b] or obtaining turn-by-turn navigation and directions. Industry has been quick to pick up on this trend; many popular LBSs are, in fact, platforms for social location sharing, including applications like BrightKite [2007], Loopt [2005], Plazes [2004], Latitude [Google, 2009], Glympse [2008], Foursquare [2009], and Places [Facebook, 2010].

Yet, despite the steady increase of social location sharing applications (LSAs), these services are still years away from mainstream adoption [ABI Research, 2008]. Past work has stated that there are at least three challenges preventing the widespread adoption of location-aware computing: 1) the lack of low-cost, convenient location finding technologies, 2) inadequate techniques to address end-user concerns about location privacy, and 3) the lack of useful, usable location-based services [Hong, 2003].

The growing ubiquity of GPS capable mobile phones at least partially addresses the first adoption barrier. With the US mobile market penetration rate at nearly 90% [CTIA, 2008], one can easily consider the mobile phone to be the preferred ubiquitous device for the everyday user. Furthermore, many of these phones support location-aware capabilities using either embedded hardware, like GPS chipsets, or through additional software protocols, like WiFi fingerprinting (e.g., [Cheng, Chawathe, LaMarca, and Krumm, 2005; LaMarca, Chawathe et al., 2005; Schilit, LaMarca et al., 2003]) or cell-tower triangulation (e.g., [Chen, Sohn et al., 2006; Varshavsky, Chen et al., 2006]). Of course, having a location-equipped phone does not always result in a perfect location-aware user experience. With today's technology, there are still many open technical challenges related to location sensing, including how to minimize power consumption and how to maximize sensing accuracy, particularly when relying on non-GPS technology. However, even with these challenges, the current state of mobile positioning technology is more or less sufficient for most, if not all, social location sharing applications. Moreover, these technical challenges can be studied independently of the other aforementioned adoption barriers, namely addressing end-user location privacy and creating useful, usable location services.

We propose that *location abstractions* can help address these two adoption barriers. We define these abstractions to be more generalized descriptions of one's location. For example, instead of describing a place using an address (e.g., "417 S. Craig St., Pittsburgh, PA 15213") or latitude-longitude coordinates (e.g., "40.444, -79.949"), we can use a *less precise* description by referring to the place's geographical properties, like its neighborhood (e.g., "Oakland, Pittsburgh"), or the place's semantic properties, like its business name (e.g., "Starbucks") or its type (e.g., "coffeeshop"). By definition, more abstracted location information should inherently provide additional privacy protection for end-users. In this thesis, we intend to *empirically* verify whether location abstractions can adequately address users' perceived concerns about their location privacy, without significantly detracting from the usefulness of social location sharing.

The remainder of this chapter covers the intricacies related to end-user location privacy and the challenges for social location sharing in particular. We make an argument that there is a specific relationship between privacy and utility (usefulness), and that only be addressing both of these issues together can we come up with a solution to potentially address the adoption problem facing current LSAs.

1.1 End-User Privacy Challenges for Location Sharing

Location-aware technologies introduce significant privacy challenges for end-users. In particular, technological advancements have helped to dramatically lower the cost of sensing, recording, and sharing large amounts of users' location data. What makes this potentially more intrusive is that location-based computations can be done in real-time and in a manner that is machine readable, searchable, and easily aggregated over time. These characteristics introduce significant privacy risks, ranging from everyday risks, such as disclosing sensitive locations to your friends and family, to extreme risks, like those relating to one's personal safety (e.g., avoiding stalkers). Past work has delved into some of these privacy concerns, which can be seen in end-user interviews about location-tracking technologies (e.g., [Barkhuus and Dey, 2003; Harper, 1995; Kaasinen, 2003]) as well as several press releases regarding potential end-user privacy violations (e.g., [Liedtke, 2007; Whalen, 1995; Zuckerberg, 2006]).

Adequately addressing users' privacy concerns is vital to the long-term success of location-aware technologies. If not dealt with, then service providers risk being outright rejected by their users (e.g., [Harper, 1995]). There have been various strategies for addressing location privacy. One way is to share locations anonymously by removing unique identifiers, such as one's username. For example, instead of saying that Alice is at 123 Main Street, the application could just say that someone is at 123 Main Street, without specifying any particular user. Another privacy-preserving mechanism is for LBSs to obscure users' location information by hiding their true location amongst other users [Beresford and Stajano, 2003; Gruteser and Grunwald, 2003]. For example, if Alice

is the only one at 123 Main Street and there are ten other people scattered about Main Street, then this obfuscation technique would opt to share a more generalize location description (Alice is at Main Street), so that Alice's true location (123 Main Street) could be hidden amongst other people's location.

While both of these strategies (anonymity and obfuscation) are indeed privacy-preserving, neither of them are applicable to the newest class of location applications that support social location sharing, i.e., location sharing within a social network. In these cases, the identity of the discloser is just as important as the location data, if not more so. Consider a scenario where a user shares her location information with others in order to provide a sense of co-presence and awareness to her friends and family. Without information about the user's actual identity, the location information becomes much less meaningful to those in her social network. In these situations, it is arguably more useful for the application to provide privacy mechanisms to ensure that potentially sensitive locations are not accidentally revealed, rather than to ensure end-user anonymity. This way, the user's location information (and her identity) is shared only with a preselected group of people that she designates, who can then socially engage with her based on her location information.

It is worth pointing out that the privacy concerns we are addressing in this thesis are specifically related to location sharing between individuals. This is different from past studies that have examined privacy concerns from an ecommerce perspective (e.g., [Ackerman, Cranor, and Reagle, 1999]). These cases typically involve scenarios where users disclose personally identifiable information (e.g., email addresses, browsing behaviors, birth dates) to businesses or government organizations. The social dynamics and the privacy expectations between a consumer and a business (or government) are vastly different than that of between two individuals (e.g., between family members or friends). In the latter case, the *social* aspect of the relationship primarily defines the information exchange, as opposed to factors like financial incentives (e.g., when sharing with ecommerce organizations) or societal duties (e.g., when sharing with the government). In this thesis, one of our goals is to better understand the privacy

implications for location applications that specifically support location sharing within one's social network. In particular, we are interested in examining end-users' *perceived* privacy concerns about social location sharing. Whereas *actual* privacy concerns may focus more on a computational analysis of the end-users' privacy (e.g., using quantitative analysis via data mining [Jones, Kumar, Pang, and Tomkins, 2007] or behavioral economics [Acquisti, 2009]), we intend to look at privacy though a more subjective lens and evaluate privacy in terms of how comfortable users may or may not feel when engaged in social location sharing.

The challenge behind evaluating subjective end-user privacy concerns is that location privacy is not a discrete phenomenon that can be described as being on or off. Instead, there are degrees of privacy and often time privacy concerns are better expressed as shades of gray. In addition, privacy concerns are malleable, susceptible to current societal norms, and can change over time due to both positive (e.g., reconnecting with friends) and negative exposures (e.g., becoming a victim of identity theft). Thus, this thesis work presents a first step in systematically understanding how users' utilize social location sharing. To do this, we combine both *quantitative studies* and *qualitative feedback* to uncover how privacy factors into users' sharing behaviors.

1.2   The Problem: Privacy vs. Utility Tradeoff

It is clear that location sharing is directly impacted by privacy concerns, which can be multi-dimensional and hard to analyze [Hong, 2005b]. These challenges are compounded by the fact that there is an implicit tradeoff between privacy and utility that makes it particularly difficult for users when privacy is discussed within the framework of social location sharing. This type of information exchange is the focus of this thesis and is embodied by location-based services that rely heavily on social network information sharing. A common scenario for this type of location sharing applications (LSAs) is described below:

> *Alice is curious about her friends' whereabouts and decides to use her*
> *mobile device to look up the current locations for all of her friends.*
> *The application displays this information very precisely on a map,*
> *where each friend's location is represented by a pushpin placed at*
> *specific geographical coordinates. Alice knows that her location*
> *information is also being shown to her friends in a similar fashion. As*
> *Alice occasionally has concerns about sharing her location*
> *information, she explicitly opts out of sharing her location information*
> *with her friends. Because of this particular privacy setting, Alice's*
> *friends are no longer able to infer her whereabouts based on the*
> *application's map-based display of everyone's current location.*

This scenario describes a common disclosure model used in many social location sharing applications, where the decision to share one's location becomes an "all-or-nothing" decision. On one hand, users can opt to disclose nothing (as Alice did) and not have their information shown on their friends' map. This choice affords complete privacy to the user, but the user also misses out on any social benefits that might have resulted if he had shared his location information with others. On the other hand, user can choose to disclose everything, which for LSAs means that a highly precise description of users' current location is shared (i.e., as was the case with Alice's friends). This choice provides more opportunities for social engagement with others, such as allowing for serendipitous encounters [Barkhuus, Brown et al., 2008] and increasing awareness between loosely connected friends [Oulasvirta, Petit, Raento, and Tiitta, 2007]. But these social benefits come at the cost of revealing potentially sensitive information, as users' locations are precisely pinpointed on a map. Thus, an "all or nothing" disclosure model ultimately forces uses to choose between whether they value their privacy more (and opt to disclose nothing) or whether they value social utility more (and opt to disclose everything).

From a service provider's perspective, users should engage in information sharing, both for the sake of maximally ensuring that social value is obtained from using the service, as well as ensuring that they obtain as much data as possible from their users. As a result,

many LSAs have default sharing preferences set so that users disclose everything. However, privacy-sensitive users may not be comfortable with such blanketed location sharing and, with only a limited disclosure model, these users will likely abandon the application altogether. In the long term, by not adequately addressing *both* privacy *and* utility, service providers run the risk of alienating users, which can make it much more difficult for LSAs to maintain an active community of users exchanging information. Without enough members using the service, new users are less likely to be attracted to using the service.

Thus, the "all or nothing" disclosure model presents an inherent dilemma for end-users. In order for users to have a chance at experiencing even a hint of social utility from using LSAs, they must decide up-front that they are willing to share a very precise description of their location with others. However, disclosing such detailed information may be above many users' privacy threshold. Yet, they must engage in this level of location sharing for at least a short time period, if they desire to at least ascertain if LSAs are worth using.

We refer to this problem as the *privacy vs. utility tradeoff*. Current implementations of social location sharing applications that use an all-or-nothing disclosure model are simply not expressive enough to provide users with the means to resolve this tradeoff in a satisfying way. In fact, the gap between these two choices effectively creates a privacy barrier for many users, preventing them from fully engaging in LSAs (see Figure 2). We posit that, in order to provide sufficient privacy mechanisms for social location sharing, LSAs should provide additional disclosure options that allow users to better balance their concerns about preserving their location privacy and their desire to engage in potential social interactions, as a result of sharing their location information. These middle ground options can help scaffold the privacy barrier created by the all-or-nothing disclosure model.

Figure 2. Visual depiction of the privacy vs. utility tradeoff. Many social location sharing applications use an all-or-nothing disclosure model, resulting in a privacy barrier that prevents many users from comfortably engaging in location sharing behaviors due to privacy concerns.

1.3   A Solution: Location Disclosure Abstractions

One intuitive solution is to supplement the all-or-nothing disclosure model with additional disclosure options that lie in between the two extremes; we refer to these options as *location disclosure abstractions*. We use the term *abstractions* to emphasize that these disclosure options are less descriptive and less precise than the full disclosure option that is usually represented as geographical coordinates on a map. The advantage of offering abstractions is clear: abstractions provide users with additional flexibility in how they would like to describe and share their location information with others. As a result, they are more likely to feel comfortable participating, at least to some degree, in location sharing behaviors. Providing abstractions is also beneficial for service providers since users who are more comfortable with location sharing are also more likely to continue using the service; thus, LSAs will obtain at least a partial description of these users' data and may be able to retain users who would have originally shied away from using LSAs. While stakeholders like service developers are bound to prefer more descriptive location information, it is likely that they can still learn useful information from less precise descriptions which can, for example, still be helpful when addressing general usability concerns. In addition, by appealing to a wider audience, a more expressive disclosure

model can potentially lead to higher application usage from the average user, as well as more sustained application use.

Past work has looked at several types of location abstractions. Two examples of these options have included: *geographic abstractions* and *semantic abstractions*. Because location data is hierarchical in nature, it lends itself well to geographic abstractions (also referred to as location blurring [Hong, 2004]). In these cases, locations can be intuitively described along a spectrum, depending on how precise of a description one wants to share with others; these abstractions can range from the user's current street address or nearest intersection to the users' current neighborhood or city to the user's current state and country. Semantic abstractions are instances when locations are described using place labels that refer to the *place* (social qualities) vs. the *space* (geographical nature) of the location [Harrison and Dourish, 1996]. Common examples of semantic abstractions used in past work (e.g., [Lin, Xiang, Hong, and Sadeh, 2010]) include referencing the type of place it is (e.g., a coffee shop, restaurant, shopping mall), a personal label (e.g., "my home", "my workplace"), or a business name (e.g., Starbucks, McDonald's).

These two types of location abstractions are arguably a much better match to how users normally describe their locations to others in daily conversations than the fully precise description based on geographical coordinates [Laurier, 2001; Weilenmann, 2003]. However, there are certainly other motivations for picking location abstractions. From a computational perspective, some examples of relatively easy abstractions include using the user's current time zone or the user's current state of motion [Bentley and Metcalf, 2007] as an indirect representation location information. In Bentley and Metcalf's study, the shared information is a binary choice between being labeled as "moving" or "not moving". When users are not moving, it is assumed that they have arrived at a place and, when they are moving, it is assumed that they are between places. This type of location abstraction provides a strong level of privacy protection in that information receivers must have a significant amount of inside knowledge to be able to resolve which specific place a user may be at. While this can be useful for addressing the privacy vs. utility tradeoff in close-knit relationships (e.g., between spouses/ significant others and

immediate family members), sharing location information that has been too strongly abstracted can result in location sharing behaviors that are much less useful for more weakly connected relationships. For these types of people, they are unable resolve ambiguities that may arise when they only have access to vague location descriptions. Thus, sharing grossly abstracted information is unlikely to be very useful in LSAs that rely on sharing within a mixed social network (i.e., one that includes both strong and weak social relationships).

Table 1 provides a summary of the pros and cons for geographic and semantic location abstractions. Most importantly, we see that there are computation advantages to using geographic abstractions, though there are potentially more social benefits to using semantic abstractions.

From a service provider perspective, the intention behind incorporating either type of abstractions is the same: it gives the user additional disclosure options so that they can better address the privacy vs. utility tradeoff and, in turn, feel more comfortable in using location sharing applications. While varying a location's description will inevitably impact the degree to which users may experience the social benefits associated with LSAs, it is assumed that, by sharing at least a partial location description, users will begin to appreciate the potential for such social benefits. Then, by exposing users to the potential upside of social location sharing, our hope is that users will be able to more accurately judge whether LSA's social utility is worth the privacy risks inherently associated with location sharing behaviors. In essence, these abstractions serve as a way for LSAs to scaffold the privacy barrier that have been traditionally associated with the all-or-nothing disclosure model, as seen in Figure 3.

	Geographic Abstractions	**Semantic Abstractions**
Pros	Computationally easy to determine • Often based on GPS readings from mobile devices • Direct consumption of GPS readings are usually sufficient to extract location information Some instances of social relevance for less descriptive geographic abstractions • In day-to-day conversations, people do reference city-level abstractions. However, other geographic abstractions are less often used.	Many examples of conversational use of semantic abstractions • In casual references to location information in day-to-day conversations, it is common to find semantic references Offers many levels of descriptiveness, which can be useful for providing degrees of privacy in terms of choosing what location information to share • Often privacy can be preserved while still providing meaningful awareness (e.g., sharing "coffee shop" as a generic semantic abstraction vs. sharing "Pittsburgh" as a city-level geographic abstraction)
Cons	Limited social relevance for highly descriptive geographic abstractions • Outside of navigation purposes, people rarely express their location in terms of specific geographic coordinates or street addresses. Can lead to end-user privacy concerns • When used in "all or nothing" disclosure models, users may feel compelled to share more than they feel comfortable with • Users are often not comfortable sharing a map of their home location [Tsai, Kelley, Cranor, and Sadeh, 2009]	Is error-prone and can be difficult to automatically compute • Accurate semantic labels are dependent on having accurate geographic coordinates (as sensed by the user's mobile device) • Translating raw GPS readings to a semantic label is highly dependent on having up-to-date databases (e.g., for retrieving business names)

Table 1. An overview of the potential pros and cons for using geographic and semantic abstractions. Geographic abstractions tends to have more computational advantages, while

**Figure 3. By adding location abstractions, we hope to lessen the privacy barriers
that users may feel when engaging in social location sharing.**

1.4   Defining the Dissertation Scope

So far, we have introduced the problem facing many LSAs, namely that users are faced
with a tradeoff between privacy vs. utility when making decisions about social location
sharing. We have also introduced a potential solution to this problem, which is to use
location abstractions to provide a way to alleviate the seemingly large privacy barrier
presented by all-or-nothing disclosure model. In this thesis, we are less interested in the
types of abstractions that could be used in LSAs. Instead, we are interested in exploring
how abstractions can be effectively used in LSAs. In other words, we make the
assumption that it is in the best interest of the service provider to design LSAs to include
the concept of location abstractions.

Furthermore, this thesis focuses on two types of abstractions: geographic location
abstractions and semantic location abstractions. There are two reasons that we focus on
these two abstraction types. First, these abstractions are meaningful to a large range of
relationship types. Past work has shown that, in most cases, a significant proportion of
online social networks include weak social ties (e.g., casual friends and acquaintances)
[Donath and boyd, 2004; Wellman, Haase, Witte, and Hampton, 2001]. Given that we are
focusing on LSAs that support similar types of social networks, it is important that we

choose an abstraction that will appeal and be meaningful to a diverse set of relationship types. Second, these two abstractions types already familiar constructs and are frequently used in day-to-day conversations [Weilenmann, 2003]. Thus, incorporating them into LSAs should not introduce any addition cognitive burden for users to understand and utilize.

A large portion of the research questions discussed in this thesis will examine how well location abstractions address end-user privacy concerns for social location sharing, in a context that is demonstrated by LSAs that support exchanging location information within a social network. But, as we have previously discussed, social location sharing is not just about privacy; it also touches upon issues relating to social utility as well. Consequently, we also intend to explore how location disclosure abstractions can help address the privacy vs. utility tradeoff commonly encountered in LSAs.

In particular, we are interested in how geographic and semantic location abstractions compare in terms of what they offer to users both privacy-wise and utility-wise. Table 1 highlights some of the practical differences between geographic and semantic location abstractions. For example, geographic abstractions are relatively easy to incorporate in location-based applications, as there is usually a direct translation between the raw GPS readings and the shared location information. For semantic abstraction, the translation process is much more complex, though past work has suggested that such abstractions are better matched in conversational sharing of location information. So, from a computational perspective, there are already important differences to consider. We hope to further probe these distinctions to determine if there are also other implications to consider when designing location abstractions for LSAs. In particular, we intend to provide a deeper understanding of the *human* perspective for these abstractions. In other words, do users have a preference for a particular kind of abstraction and what reasons are there for their preference?

Throughout this thesis, we are interested in providing *empirical evidence* for how location abstractions can both positively and negatively impact social location sharing.

To study this issue, we look at four different areas that location abstractions can impact LSAs. First, we examine how location abstractions affects users *make decisions about location sharing* and whether it impacts people's perceptions of location privacy. Second, we look at how location abstractions affect how users *configure their privacy settings* for social location sharing applications. Third, we explore how to leverage location abstraction in various information visualizations and whether different ways of *presenting location information* affects people's perceived privacy concerns. Lastly, we examine potential *outcomes for location sharing* and pay particular attention to metrics for capturing *actual privacy risks* and indirectly measures for *social interaction*.

These four research areas span both privacy and social utility concerns. Our thesis is that location abstractions (i.e., geographic and semantic abstractions) can help address these concerns from both the users' perspective (in terms of *perceived* privacy concerns and utility), but also *actual* privacy concerns and utility too. Specifically, our thesis is as follows:

> *By providing geographic and semantic disclosure abstractions, social location sharing application can better address end-user privacy concerns in at least three ways. First, abstractions can simplify privacy configurations. Second, abstractions provide more visualization opportunities that can be both engaging and privacy-sensitive. Third, using abstractions can lead to higher self-reported privacy comfort levels, while also providing some degree of social utility, as measured by online social interaction patterns.*

1.5   Dissertation Contributions

Throughout this dissertation, there are a number of major contributions:

- A framework that distinctly identifies social location sharing as being separate and unique from past studies of location sharing. (Chapter 4)

- Empirical evidence that users have different reasoning processes for social-driven vs. purpose-driven location sharing, resulting in different sharing preferences and different privacy preferences when engaged in location sharing scenarios. (Chapter 4)

- Quantitative evaluations for a month-long field deployment of a location-aware social application. (Chapter 5)

- Empirical evidence that users of a social location sharing application are comfortable with a disclosure protocol that includes geographic location abstractions and prefer these location abstractions as their default sharing policy. (Chapter 5)

- Empirical evidence that, when compared to an all-or-nothing disclosure model, location abstractions can result in users sharing more location information with more relationship types. (Chapter 6)

- Empirical evidence that, when compared to an all-or-nothing disclosure model, location abstractions can result in simpler rule-based privacy configurations. (Chapter 6)

- Descriptive analysis of different visualizations for sharing location history. (Chapter 7)

- Analysis for how varying visual presentation of location information can impact users' perceived privacy comfort levels, leading to design suggestions for how to design privacy-sensitive data visualizations for location trails. (Chapter 7)

- Data analysis for measuring the actual privacy preservation of using location abstractions. (Chapter 8)

- Data analysis for indirectly measuring the social interaction opportunities for sharing location abstractions within an online social network application. (Chapter 8)

- Various user study methodologies to explore perceived and actual privacy concerns using a controlled lab environment (Chapters 4, 6, and 7), a field deployment (Chapter 5), and only a data-centric analysis (Chapter 8).

- A set of empirical studies that systematically evaluates end-user privacy concerns for social location sharing for a broad range of relationships, ranging from close

social ties (e.g., immediate family members and significant others) to weak social
ties (e.g., causal friends and acquaintances).

- A set of empirical studies that examines social location sharing issues from
several perspectives: both asynchronously (Chapters 4, 6, and 7) & and
synchronously (Chapter 5), and sharing current (Chapter 5) and past location
information (Chapters 4, 6, and 7).

1.6   Dissertation Outline

The remainder of this dissertation is organized as follows:

Chapter 2 provides a discussion about past work in location sharing applications and, in
particular, studies that evaluate end-user privacy concerns. In Chapter 3, we provide
several different ways of framing the related literature. In this chapter, we also discuss
our research questions and describe a framework to structure the results of this
dissertation work.

Chapters 4 through 8 provide details about the five user studies that constitute this
dissertation work. In Chapter 4, we examine how users make decisions in regards to
social location sharing, as opposed to other types of location sharing. In Chapters 5 and 6
examine the impact that location disclosure abstractions have on end-user decision
making and their privacy configurations. Chapter 7 describes a study evaluating how
specific types of information visualizations (all of which make use of location
abstractions) can impact users' location sharing preferences. Chapter 8 presents results
from analysis investigating the social benefits of sharing location abstractions, as opposed
to fully precise descriptions of one's location information.

In Chapter 9, we synthesize the results from the five user studies that encompass this
thesis. We also discuss possible design implications for future social location sharing
applications, paying particular attention to the design of privacy-related features.

Chapter 10 concludes the dissertation by discussing the limitations of our work and topics for future work.

2. Background: Privacy & Location Sharing

In ubiquitous computing systems, privacy is often relegated to the sidelines, not because researchers don't acknowledge its importance but because privacy is difficult to describe, analyze, and assess [Hong, 2005b]. In addition to technical considerations, privacy touches upon legal issues, corporate policies, and societal norms. Given people's tendencies for bounded rationality [Acquisti, 2005], it is understandable to find that users often have difficulties reconciling all these dimensions when quantifying their privacy preferences.

Evaluating privacy concerns can also be challenging since privacy can frequently seem like a moving target, either changing gradually over time (e.g., due to increasing exposure to new societal norms) or changing very quickly (e.g., due to recent negative experiences in privacy loss). Because of privacy's malleable nature, users often differ in their individual perceptions about privacy, their priorities regarding privacy concerns, and their reasons when making privacy-related decisions.

Privacy also suffers from being extremely context-dependent. A user's rationale for making certain privacy decisions in one domain may not transfer to other domains. Consider the public's adverse reaction [Lynch, 2007; Zuckerberg, 2007] to Facebook's launch of the Beacon service [Facebook, 2007],which allowed Facebook to track users' purchasing behaviors on third-party websites like Amazon, Barnes & Noble, etc. Many people found that it was inappropriate for Facebook to share this information with others in their online social network. In contrast, consider how many people routinely use credit cards and loyalty cards when shopping. These consumers normally report that the convenience and potential monetary savings (often very small) outweigh the privacy cost

of having their financial transactions tracked by businesses (and potentially being misused as a result) [Acquisti, 2004]. Though the type of information being tracked is very similar (i.e., both are recording ecommerce-related activity), these two scenarios solicit very different user reactions, demonstrating that it is important to consider the *task* and *context* in which users are making decisions about their privacy.

In this dissertation, we confine our discussion of privacy to the domain of location sharing. In the following sections, we provide an overview of past location-aware applications that have featured location sharing between individuals. We also describe past work that is related to evaluating end-user privacy concerns for these types of applications.

2.1   Overview of the Different Kinds of Location Applications

Smith et al. noted that an emerging class of pervasive computing are applications that "share location information in social communication" and referred to these as social location disclosure applications [Smith, Consolvo et al., 2005b]. This dissertation primarily focuses on this type of location-based service, though we have rephrased this to be *social location sharing application.* In particular, we are only considering the subset of location-based services that support *social sharing* of location information *within a social network* and in a *non-anonymous* manner.

These additional constraints eliminate three types of location services which do not fall within scope of this dissertation. First, we do not focus on applications where location information is publicly broadcasted and meant to be viewed by everyone. For example, though Geonotes [Espinoza, Persson et al., 2001], E-Graffitti [Burrell and Gay, 2002], Sharescape [Ludford, Priedhorsky, Reily, and Terveen, 2007; Reily, Ludford, and Terveen, 2008], and Microsoft's SlamXR [Counts and Smith, 2007] all have operating modes that support private location sharing (i.e., where users target their location sharing

to specific individuals), these community-oriented applications primarily encourage users to make their location information publicly accessible.

Another class of location-aware application that we do not consider are those that anonymously broadcast location information (e.g., Hitchhiking [K.P. Tang, Keyani, Fogarty, and Hong, 2006]). In this dissertation, we are primarily concerned with evaluating privacy concerns for social location sharing. Anonymous information sharing, by definition, neutralizes many privacy concerns, as there is no explicit link between the discloser's identity and her location. However, anonymous sharing cannot support many social scenarios, such as providing social awareness of your friends' current whereabouts. In these cases, without any identity information, one cannot know for sure who is at any particular location (only that some person is at a location). Because of these features, the privacy questions that we are examining are irrelevant for location applications that support anonymous location sharing.

Finally, we also do not consider location applications that only use location information for personal informatic purposes. Systems that fall in this category tend to track location information in combination with other contextual data (e.g., fitness-oriented applications like RunKeeper [2008]), for personal planning purposes (e.g., travel applications like Dopplr [2009]), or for information retrieval purposes (e.g., location-based search engines like Where [2009] and Yelp [2004]).

The intention behind excluding these particular types of location applications is so that we can instead focus on services that broadcast location information within a pre-defined social network, for the purposes of enhancing social awareness. We refer to such services as *social location sharing applications*, which, as previously mentioned, is simply a subset of Smith et al.'s definition for social location disclosure applications.

2.2 Organizational Frameworks for Social Location Sharing Apps

In this section, we present three different ways for framing existing social location sharing applications. The first way is to classify applications according to how many people are expected to receive the location information. More specifically, we can classify applications based on how large of a social network the application was originally intended for. Many social LSAs target relatively small groups with relatively homogenous social relationships, e.g. a group of co-workers, immediate family members, or a close-knit group of friends. Examples of such systems include SLAM [Microsoft Research, 2009], Radar.net [2009], and PlaceMail [Ludford, Frankowski, Reily, Wilms, and Terveen, 2006]. A few of these applications can be scaled up to larger groups, such as Twitter's geolocation status updates [2009a], dodgeball [2009], SWARM [Farnham and Keyani, 2006; Keyani and Farnham, 2005], and ContextContacts [Oulasvirta, Raento, and Tiitta, 2005]. In contrast, there are only a few social location sharing applications designed for very large populations and are capable of supporting location sharing for tens of thousands of users. Most of these applications fall outside the scope of this thesis because they tend to either anonymously broadcast users' location information (e.g., Hitchhiking [K.P. Tang, Keyani, Fogarty, and Hong, 2006]) or they publicly broadcast users' location information (e.g., GeoNotes [Espinoza, Persson et al., 2001]).

A second way of organizing social location sharing applications is by examining how the location information is shared between users. For example, we can classify LSAs according to whether the location disclosure occurs synchronously or asynchronously. Examples of synchronous applications include People Finder (now called Locaccino) [Cornwell, Fette et al., 2007; Sadeh, Hong et al., 2009], Reno [Smith, Consolvo et al., 2005a], Motorola's motion presence application [Bentley and Metcalf, 2007], Awarenex [J.C. Tang, Yankelovich et al., 2001], and WatchMe [Marmasse, Schmandt, and Spectre, 2004]. For the most part, these applications tend to only provide *nearly* real-time location awareness, so one might technically consider these LSAs only weakly synchronous, particular when compared to mediums like voice communication or video conferencing (which both support more strict interpretations of real-time data exchanges). However,

for comparative purposes, these LSAs sufficiently demonstrate synchronous location sharing.

Asynchronous sharing, on the other hand, emphasizes ad-hoc location awareness. Examples of such location services include DeDe [Jung, Persson, and Blom, 2005], PlaceMail [Ludford, Frankowski, Reily, Wilms, and Terveen, 2006], Groovr [2009], and comMotion [Marmasse and Schmandt, 2000]. In most of these applications, when location information is shared with others, it is often to indicate that the user was *previously* or *recently* at that place, rather than that she is *currently* at that place. It is important to note that both synchronous and asynchronous location-aware social applications support social awareness. It is only the *type* of awareness that differs: one provides real-time location updates (synchronous awareness) and the other provides a history (complete, partial, or otherwise) of past locations (asynchronous awareness).

A third way of organizing social location sharing applications is by how the application delivers location information to the user. Traditionally, location-based services share data using either a push- or pull-based model. Pull-based location services provide on-demand access to location information. For example, in AT&T's FamilyMap [2009], the user is provided with a buddy list and, in order for location information to be exchanged, a user must click on a username in her buddy list in order to see that person's location plotted on a map. On the other hand, push-based location services provide location information continuously to the user. Google Latitude [2009] and Loopt [2005] both support this type of location sharing and are currently implemented by having a continuously updated map showing the most recent location for a specific list of users.

Figure 4 provides a visual representation for two of the three organizational schemes that we just discussed. The diagram categorizes social location sharing applications according to scale (i.e., how many people was the application optimally designed for) and information delivery (i.e., how is information shared and/or exchanged with others). By examining where these LSAs overlap, we can provide a much more precise definition of the type of LSA that we intend to focus on in this dissertation. In particular, when

referring to social location sharing applications, we are specifically referring to applications that:

- share location information *for social awareness* purposes (asynchronously or synchronously)

- share location information *within a social network* (*medium-scale*)

- share location information using a *push-based* model (information delivery)
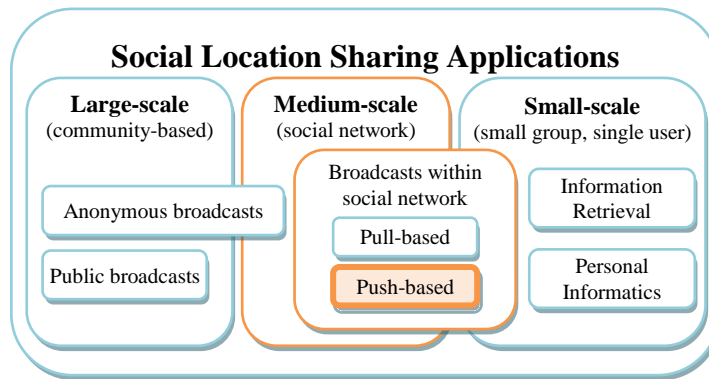
Figure 4. Classification of social location sharing applications according to scale (large, medium, or small) and the type of information delivery (push- or pull-based, personal consumption, or search/retrieval purposes). The orange highlighted box (i.e., medium-scale, push-based applications that broadcast location information within a social network) defines the type of location application that this dissertation focuses on.

2.3   A Framework for Examining Location Sharing Privacy Concerns

There has been a fair amount of past work dealing with end-user privacy concerns about location disclosures and, more broadly, about privacy in ubiquitous computing environments. This includes theoretical frameworks for modeling how users reason about location privacy (e.g., [Consolvo, Smith et al., 2005; Hong, 2005a; Iachello and Hong, 2007; Khalil and Connelly, 2006; Lederer, Hong, Dey, and Landay, 2005]), published experiences regarding location privacy following a deployment of a ubiquitous computing system (e.g., [Harper, 1995; Hindus, Mainwaring, Leduc, Hagström, and

Bayley, 2001; Hindus and Schmandt, 1992; Hong, 2005a; Kaasinen, 2003]), and firsthand descriptions of users' experiences about location privacy (e.g., [Hong, 2005b; Lederer, Hong, Dey, and Landay, 2005]).

Starting with work by Bellotti and Sellen [1993], there has been a general consensus that providing adequate controls and feedback mechanisms is essential for users to successfully manage their information privacy. Consequently, most privacy-related studies about location disclosures have framed their discussions around these two parts of the privacy "equation" (controls and feedback), though there is more work exploring issues relating to the design and evaluation of privacy controls for ubicomp systems.

In this section, we use a different framework to describe past privacy work. Specifically, we examine end-user privacy concerns from a timeline perspective and look at these issues at three different stages of a typical location exchange. Table 2 describes this disclosure timeline by outlining privacy concerns that typically happen before location information is exchanged (this is where much of the past work on privacy controls would fall under), during the exchange of location information, and after location has been exchanged with others (this is where much of the past work on privacy feedback mechanisms would fall under). Thus, our framework expands upon Belotti and Sellen's work by including a third category of privacy-related research issues that are important to consider for location-aware ubicomp systems.

	Stage 1: Before disclosure	Stage 2: During disclosure	Stage 3: After disclosure
Privacy concerns	Who sees what data? When to share data?	How is the data shared?	What data was shared? Who has seen the data?
Privacy features	Privacy controls: settings & configuration	Disclosure protocols & interaction styles	Privacy feedback: access logs & notifications

Table 2. A framework showing privacy concerns about location sharing, described in relation to a disclosure timeline (before, during, and after location disclosures). The privacy features, listed in the second row, thematically describe how past related work fits in this framework.

To better understand the disclosure timeline, imagine a user who is just starting to use a typical social LSA. When the user decides to engage in location sharing behaviors, she will first configure her sharing preferences using the application's privacy controls (Table 2, "Stage 1: Before disclosure"). This can be done for each location disclosure or by initially configuring a default privacy policy that will be applied to all of her subsequent location disclosures with others. When exchanging location information, how the information is shared (e.g., is reciprocity enforced, how is the information visualized, etc.) become important privacy issues to consider (Table 2, "Stage 2: During disclosure"). After exchange location information, applications can decide whether to notify users and can vary how much feedback they provide about how much location information has been shared on the user's behalf (Table 2, "Stage 3: After disclosure").

In the next section, we will use this framework to review past literature that relates to studying end-user privacy concerns for social location sharing.

2.4   Overview of Privacy Mechanisms Used in Location Sharing

The first stage of the disclosure timeline occurs prior to sharing any type of location information with others. During this time, the user makes various decisions relating to how she will specify her sharing preferences using the application's privacy configuration interface. The application will then either apply these settings to just the current disclosure or to all future location disclosures, ensuring that any location information that is shared with others is only exchanged according to the user's privacy configuration.

When deciding what location information to share with others, past work has found that users are mostly concerned with two factors; they want to know who is asking about their location information and they want to know what is the context for why this person is requesting their location information [Adams, 2000; Brown, Taylor et al., 2007; Iachello,

Smith et al., 2005; Khalil and Connelly, 2006; Lederer, Hong, Dey, and Landay, 2005; Sadeh, Hong et al., 2008; Smith, Consolvo et al., 2005a]. Of these two factors, the requestor's identity has been found to be more important than knowing the context in which the requestor is asking for the information [Consolvo, Smith et al., 2005; Lederer, Mankoff, and Dey, 2003]. In terms of designing privacy controls for location disclosures, there have been several approaches, including group-based controls [Patil and Lai, 2005], proximity-based controls [Hull, Kumar et al., 2004], place-based controls [Sadeh, Hong et al., 2008], time-based controls (using day of week and time of day to specify sharing preferences) [Sadeh, Hong et al., 2008], controls using heuristics (like those using case-based statistical reasoning [Sadeh, Hong et al., 2008]), persona-based controls [Lederer, Hong, Dey, and Landay, 2005], and policy-based controls [Langheinrich, 2002]. The range and diversity of these controls suggest that there is not yet a good solution for designing privacy controls for location sharing. Furthermore, many of these controls have only been evaluated qualitatively, in terms of whether participants like or dislike them; it is rare to find studies based on field deployments that have more quantitative evaluations.

In the second stage of the disclosure timeline, variations in how location information is shared can significantly impact end-user privacy concerns in at least three ways. First, depending on how the user's location is being computed, users can feel less in control of how their location information is being shared. For example, location information could be manually provided by the user or it could be sensed automatically through positioning technology embedded in their mobile devices. Automatic sensing often elicits fears of being continuously tracked by others [Gruteser and Liu, 2004] and can significantly increase the privacy burden for end-users, as they now have to worry about whether or not they have control over their location information after it has been sensed and whether they can manipulate their location information before it is shared with others. While the term "manipulation" may imply deception, it can also refer to simply changing the level of location granularity (e.g., "5000 Forbes Ave, Pittsburgh, PA 15213" becomes "Pittsburgh, PA").

Work by Lederer [2003] and Hong [2004] have both indicated the importance of providing options to obfuscate disclosures by *varying the location granularity*. In Lederer's location-aware application, he used geographic abstractions to provide four levels of granularity for his privacy controls: 1) precise (e.g., "Starbucks Café at 123 New Montgomery"), 2) approximate (e.g., "San Francisco Financial District"), 3) vague (e.g., "San Francisco"), and 4) undisclosed (e.g., "unknown"). Several past work have also suggested that varying the precision for location disclosures helps to provides users with "plausible deniability" [Harper, 1995; Hong, 2004; Lederer, Hong, Dey, and Landay, 2004; Lederer, Hong, Dey, and Landay, 2005]. In other words, more general location descriptions affords a user the possibility to more comfortably deny their (current or past) whereabouts, instead of outright lying or refusing to share their information. Several past studies of computer-mediated social relationships have also found that plausible deniability is important to support for end-users (e.g., [K. Aoki and Downes, 2003; P.M. Aoki and Woodruff, 2005; Nardi, Whittaker, and Bradner, 2000]).

Other location applications that allow users to vary one's location precision (for sharing with others) include Reno [Smith, Consolvo et al., 2005a] and the Whereabouts Clock [Brown, Taylor et al., 2007], both of which use semantic abstractions (via personal labels like "home", "work") to obfuscate a user's location. Broadly speaking, providing semantic place labels conceptually offers a more meaningful interpretation of location, as demonstrated by Harrison's work describing the differences between space (i.e., geographical coordinates describing a location) and place (i.e., a more social interpretation of a location) [1996].

The second stage of the disclosure timeline also highlights a second privacy concern related to how an application handles incoming location disclosure requests. Grudin and Horvitz [2003] presented three different *interaction styles* for managing information disclosures: pessimistic, optimistic, and mixed initiative. An application that employs a pessimistic interaction style requires users to provide their privacy configuration settings upfront, before any information is exchanged. While this requires more effort from the user initially, it theoretically affords more privacy protection, as subsequent location

disclosures will always follow the user's specified privacy settings. However, a possible disadvantage of this interaction style is that users may be overly conservative when initially providing their privacy settings, as they may over-estimate their privacy concerns since they are not familiar with the location exchange process and its social utility.

On the other hand, using an optimistic interaction style removes the burden associated with an upfront privacy configuration and instead suggests that users can, more or less, cope with an application's default settings and, in the event that it is inadequate, the user can simply re-adjust their settings on a case-by-case basis. In other words, with an optimistic interaction style, users will only revisit their privacy configuration if a disclosure mishap occurs and the assumption is that such mishaps will rarely occur. While this method requires less work from the user initially (recall that the user simply uses the application's default privacy settings), it does require the user to be fully aware of all of their disclosures after the fact, so that they can determine when something has been incorrectly shared and their privacy settings need to be changed. Thus, the cognitive effort is offloaded from the beginning of the disclosure timeline (i.e., when the user is configuring their privacy settings) to the end of the disclosure timeline (i.e., when the user is reviewing their past disclosures).

The third interaction style (the mixed initiative approach) that Grudin and Horvitz propose is meant to be a compromise between the optimistic and pessimistic approaches. This approach says that users will be interrupted each time there is a request for their location information, allowing users to have fairly tight control over how precise they want their location to be and also to whom they want to share their location information with. This type of privacy control is similar to the reactive access control mechanism used in the Grey system [Bauer, Cranor, Reeder, Reiter, and Vaniea, 2008]. The disadvantage of this approach is the potential for excessive interruptions and is clearly not ideal for contexts where location information may be frequently shared between users. Some systems have tried to address this problem using timed leases [Glympse, 2009; Lederer, Hong et al., 2003], where an application automatically shares a user's location information for a pre-determined window of time. Then, when the current time

falls outside this preset window, all subsequent location requests are denied and no location is shared.

In addition to varying disclosure granularity and exploring interaction styles, the second stage of the disclosure timeline also introduces privacy issues relating to *disclosure protocols*. For example, Jiang [2002] introduced the concept of "information asymmetry", where users only exchange the minimum amount of information necessary. In this way, users avoid over-sharing and lessen the chance for accidentally sharing sensitive information that is not useful to the receiver. Another variation that can be included in disclosure protocols is the concept of reciprocity [Bellotti and Sellen, 1993; Treu, Fuchs, and Dargatz, 2007], where both the discloser and the asker must share their location information with each other. In this case, both users mutually expose themselves to the privacy risks associated with location sharing.

It is also important to consider how much location data will be exchanged during each disclosure request. Most location sharing applications only share a single instance of location information, which is typically representative of users sending a "where are you now?" location request. In this type of disclosure model, LSAs typically send a single instance of the user's location, most often her most recent location. However, there are also applications that disclose location trails (e.g., Microsoft's SlamXR [Counts and Smith, 2007]). In these LSAs, the user's past N (where N can range from one hour to as long as one month, depending on the application) location instances are shared in response to each disclosure request. To our knowledge, there has been no privacy evaluations conducted for social location sharing applications that disclose a user's location trails to others. For LSAs that share only current location information, there have been a handful of privacy evaluations, though only a select few have been based on field deployments (e.g., [Tsai, Kelley et al., 2009]).

In the last stage of the disclosure timeline, the main end-user privacy concern is related to how much awareness is provided back to the user about their sharing history. Similar to the importance of having good controls, feedback has also been shown to be a helpful

privacy feature in ubiquitous computing systems [Bellotti and Sellen, 1993]. Specific examples of privacy feedback used in past location-aware applications include providing real-time notifications [Hsieh, Tang, Low, and Hong, 2007], access logs [Hsieh, Tang, Low, and Hong, 2007; Tsai, Kelley et al., 2009], social translucency [Erikson, Smith et al., 1999; Nguyen and Mynatt, 2002], and auditing [Tsai, Kelley et al., 2009].

Some commercial location sharing applications have opted to partially hide users' privacy feedback. For example, in some application users cannot see who has asked for their location information in the past or how much of their location information has already been shared. Both Facebook [2004a] and Twitter [2006] allow open browsing of users' status updates and, depending on the user's privacy settings, these updates may include location information (either through Facebook's Places feature [Facebook, 2010] or through Twitter's geolocation tags [Twitter, 2009b]). In both systems, users are not able to find out who has viewed their status updates and, consequently, users do not has or has not seen their location information. While such opaqueness may lend well for social browsing of other people's information [Lampe, Ellison, and Steinfield, 2006], it can also exacerbate end-user privacy concerns when users consider that contextual information like their location is being seen by more people than they imagine.

In summary, we have used the disclosure timeline (Table 2) to frame our discussion of related work by classifying past studies according to three different stages that occur when location information is shared between users. We pay particular attention to the end-user privacy concerns for each of these three parts of the timeline and note that most of the work done in this domain has traditionally focused on the design and evaluation of privacy controls and feedback mechanisms. However, when considering social location sharing, we posit that it is a more complete framework to think of the space using a process-based perspective that goes beyond privacy controls and feedback. In particular, making design decisions in regards to the disclosure protocol (e.g., how many granularity levels to offer, which interaction disclosure style to choose, deciding how much location information to share per disclosure request, etc.) can have important implications for end-user privacy concerns.

2.5   The Social Value of Location Sharing

Up until this point, the unstated assumption has been that by engaging in location sharing within a social network, users are afforded some social benefit. While there has not yet been empirical evidence supporting this claim, past literature has provided several pieces of qualitative evidence that, when considered altogether, suggests that there is indeed some degree of social utility for encouraging location sharing behaviors.

We start first with the understanding that users have already been shown to be receptive to the idea of location sharing. In particular, past work has shown that users often approach location requests very pragmatically and are willing to share their locations as long as there is a reasonable justification for the request [Consolvo, Smith et al., 2005; Khalil and Connelly, 2006]. Many of these studies were conducted as diary studies [Zhou, Ludford, Frankowski, and Tervee, 2005] or small laboratory experiments [Anthony, Kotz, and Henderson, 2007; Barkhuus, 2004; Cornwell, Fette et al., 2007; Lin, Xiang, Hong, and Sadeh, 2010; Patil and Lai, 2005], though a few have been deployments involving small pre-existing social groups [Barkhuus, Brown et al., 2008; Smith, Consolvo et al., 2005a].

Next, we see that past work has also shown that engaging in location sharing can help increase one's social awareness of others. An ESM study by Anthony et al [2007] found that many users disclosed their location as an indication to their friends that they were socially available. This study suggests location sharing can indirectly provide awareness of one's activity and availability for social interactions. These examples of *social* awareness are in addition to other, more commonly touted benefits of location sharing, including using location reports for "okayness checking" (e.g., making sure that a plan is on-track or making sure someone has arrived at home safely) [Iachello, Smith et al., 2005], micro-coordination (e.g., arranging, on the fly, to meet someone at a preset location) [Colbert, 2001], and coarse-grained coordination (e.g., assessing whether it is a good time to call someone) [Oulasvirta, Raento, and Tiitta, 2005]. Certainly, social location sharing does not preclude these more utilitarian purposes. However, with social

sharing, the utility is often less clear to users, so it is important to underline how past work has found that this type of location sharing can be helpful in maintaining relationships through increased awareness.

Lastly, past work has also shown that increased social awareness can often lead to better social interactions. This association has been suggested because ambient awareness can provide helpful presence information as well as help support more socially-oriented goals. For example, Nagel's Family Intercom used context to infer one's availability [2001], Avrahami et al. used context to infer one's interruptibility [2007], Awarenex [J.C. Tang, Yankelovich et al., 2001] used context to aid general communication and coordination efforts among distributed or highly mobile workers, and Bentley et al's [2007] motion-based LSA found that location sharing helped users infer others' statuses during daily routines. We posit that location awareness is just one facet of contextual awareness and can therefore be helpful in informing information requesters about a users' current status with relatively low overhead costs. Thus, adding location information can be useful for streamlining information sharing to be done at more opportune times (i.e., when one's more available, more interruptible, and open to communication), which can lead to "better" social interactions in the sense that users will be arguably more attentive in the information exchange when the user is engaged at appropriate (non-busy) times.

In terms of supporting socially-oriented goals, many past systems that have shown that ambient awareness can be helpful in terms of supporting social dynamics for groups. For example, Babble [Erikson, Smith et al., 1999] incorporated a "social translucency" feature that showed how much each user was engaged in information sharing with others. Based on a field deployment of Babble, results showed that information sharing (via social translucency) helped provide general awareness of others' social activity, improved social cohesion with others, could be used to apply peer pressure to others (to also share their information), and helped groups conform to social conventions (relating to how much information they should share). In Connecto [Barkhuus, Brown et al., 2008], a mobile micro-blogging system where users shared their location information plus a custom status message, participants often used location information as a starting

point for discussion and for ongoing play. This result suggests that social location sharing can help not only increase ambient awareness, but also help generate discussions and conversations within a social network.

By daisy chaining results from past user studies, we see that: 1) under certain circumstances, users are open to location sharing, 2) location sharing leads to improved social awareness, and 3) social awareness leads to better social interactions with others. Thus, we can indirectly posit that location sharing can, at least indirectly, provide social benefits for users. Anecdotally, location sharing has also been linked to benefits like conversational grounding (i.e., using location information as a starting point for later conversations) and serendipitous interactions (e.g., seeing that a friend you haven't seen in a long time happens to be nearby).

Furthermore, these types of social benefits are arguably more meaningful when considering location applications that target sharing within medium to large-sized social networks. In most past studies, location sharing has been explored in relatively small groups, like Connecto's study with a small group of 5-6 close friends [Barkhuus, Brown et al., 2008]. In larger social networks (similar to those found in online social network sites like Facebook [2004a]), sharing location information may raise additional privacy concerns for users since potentially sensitive information could be shared with a much more diverse group of people. Online social networks typically include several different relationship types, ranging from close social connections, like family members and close friends, to relatively weaker social connections, like casual acquaintances and professional contacts. In fact, a large proportion of online social networks are often characterized as having weak social ties to a user [Donath and boyd, 2004; Wellman, Haase, Witte, and Hampton, 2001]. For these relationship types, social bonding is almost exclusively supported using computer-mediated communication tools (e.g., IM and email), and not through face-to-face interactions or phone calls (as would most likely be the case for immediate family members and close friends). In these cases, we posit that sharing context information, like location, can provide a relatively low-cost outlet for

information sharing that can potentially help bridge the social awareness gaps when users are not actively using computer-mediated communication tools.

We have also seen examples of social location sharing in commercial applications like Foursquare [Foursquare, 2009].With more than half a million users and 15.5 million check-ins [Parr, 2010], Foursquare has generated a significant amount of user activity around social location sharing. In particular, 77.4% of Foursquare users have posted at least 30 check-ins in a month; 79.2% have checked into at least 25 different places; 57.4% have checked into at least 50 different places; and 27.5% have checked into 10+ places in a twelve-hour period at least once [Foursquare Grader, 2009]. In May 2010, Foursquare reported that users where checking into 10+ places per second [Van Grove, 2010]. Given the size of Foursquare's network, one could arguably claim that these user statistics are proof enough that social location sharing has some intrinsic social value for users.

2.6   Summary

In this chapter, we provided an overview of location-based services, paying particular attention to the types of *location sharing applications* that we will discuss in this dissertation. We also presented a three-stage framework for examining privacy-related concerns regarding location sharing, according to a timeline for how location disclosures take place. We used this framework to structure our discussion of related work in terms of privacy controls (i.e., what happens before information is disclosed), disclosure protocols (i.e., what happens while information is being shared), and privacy feedback mechanisms (i.e., what happens after information has been disclosed). We then examined literature to explain why there is a social benefit for sharing location with others in your social network. We found several examples from past literature that suggests that sharing contextual information, like location, can provide social awareness and that this type of awareness can help strengthen social relationships. These findings form the basis of our motivation for exploring abstractions in social location sharing applications. In other

words, given that there is a social value in location sharing, we are interested in seeing whether sharing location abstractions (as opposed to more precise location descriptions) can better address users' privacy concerns so that users can more comfortably partake in the benefits of social awareness.

3. Defining the Research Questions

In previous chapters, we have described privacy as a multifaceted problem that goes beyond just providing secrecy for users (e.g., through anonymity or encrypted information exchanges). In particular, this dissertation focuses on understanding privacy from the perspective of end-user comfort levels and social utility. We highlight these two dimensions because, without adequate privacy features, LSAs cannot sufficiently support plausible deniability, which prior work has shown is particularly important for location sharing. Most prior privacy studies have explored location sharing in terms of how costly it is to share that information. We argue that for *social* location sharing, the key challenge is in designing a disclosure protocol that addresses the privacy vs. utility tradeoff, which we previously described in Chapter 1.2 as being a significant barrier for mainstream LSA adoption. Our ultimate goal is to design LSAs in such a way that supports *both* end-user privacy concerns *and* allows them to appreciate at least some of the social benefits of location sharing.

3.1   Location Disclosure Abstractions

The most appealing characteristic of using disclosure abstractions is its simplicity. Conceptually, location abstractions are an extension of prior work that has already underlined the importance of supporting varying levels of location granularity in order to address privacy concerns. In our dissertation, we have also chosen to focus on two specific types of abstractions (i.e., geographic abstractions and semantic abstractions) that have already been introduced, at least to some degree, in previous instances of LSAs.

For instance, several past work have incorporated geographic abstractions in their location sharing applications [Consolvo, Smith et al., 2005; Hong, 2004; Lederer, Hong et al., 2003]. In these cases, the most detailed location description is usually a street address or a geographic coordinate pair. By using the geographic abstractions, these LSA are able to blur the precision from, for example, "Forbes Ave & Morewood Ave, Pittsburgh, PA 15213" (a street address description) or "40.443444,-79.943819" (a latitude longitude coordinate description) to "Pittsburgh, PA" (a city-level geographic abstraction). Some commercial location-aware systems have also embraced this type of abstraction. For example, Google Latitude [2009] provides three disclosure options for its users: disclosing no location, disclosing a fully precise location (i.e., the equivalent precision of a latitude-longitude coordinate description), or disclosing only a city-level location label.

In Chapter 1.3, we gave an overview of several types of semantic abstractions that have been used in past location sharing systems, including using motion (e.g., "moving" or "not moving" [Bentley and Metcalf, 2007]) and using personal labels (e.g., "home", "work"). In some systems, these personal labels are pre-determined and is the same for all users, as in the Whereabouts Clock [Brown, Taylor et al., 2007]. In systems like Reno [Iachello, Smith et al., 2005], personal labels are manually created, either initially when configuring the system or when the user arrives at a particular location (i.e., event-based).

In this dissertation, we assume a slightly different implementation for *semantic abstractions*. The Whereabouts Clock [Brown, Taylor et al., 2007] uses a small, finite set of location labels ("home", "work", and "school"), but for social location sharing such a small set of labels may be too restrictive. A typical user is likely to visit more than just "home", "school", and "work" in their daily routines. For example, they may visit places like their favorite coffee shop, their local library, or even their local grocery store. But relying completely on the user to always manually provide a label (as in the case with Reno [Iachello, Smith et al., 2005]) can be tedious and potentially disruptive for the user since they have to fairly attentive in order to faithfully completely that task for each place they visit.

Instead, in our work, we automatically generate semantic labels by querying publicly available databases or web services, such as Microsoft's MapPoint [2000], Google Maps [2005], and Wikipedia [2001a]. This method has both advantages and drawbacks. On one hand, we avoid having to interrupt the user to label each location that needs to be shared; however, the accuracy of the automatically generated label depends heavily on the quality of the database. A more detailed description of how we generate semantic labels is given in Chapter 7.6, including an in depth discussion of its limitations.

3.2   Exploring Important Privacy-Related Usability Issues

When evaluating location abstractions, past work has mostly focused on one type of abstraction (geographic abstractions) and how privacy controls can make better use of this in their designs. However, even in these studies, abstractions are typically examined only in the context of a lab setting and not through any significant field deployment. Furthermore, there has not yet been any work done to consider other types of usability issues, beyond those relating to the UI design of privacy controls, when discussing disclosure abstractions for location sharing. To address this oversight, this dissertation takes an end-to-end perspective on examining the practical implications of incorporating abstractions into LSAs. While we also touch about topics relating to privacy controls, we go much further and also look at how users make privacy decisions, what factors influence users' preferences for location visualizations, and what types of outcomes can be expected from adding these abstractions to LSAs. In the remainder of this section, we go each of the four research topics in more detail.

Our first research question focuses on how users *reason* about location sharing. Past studies of location sharing have almost exclusively considered scenarios where location is shared for more functional purposes (e.g., for collaboration or coordination), but, as we have described in previous chapters, *social* location sharing presents challenges that are unique from many past work. This dissertation will look at how social location sharing differs from other types of sharing and whether users make different decisions (and why)

about what location information to share, depending on the context of the location request.

The second research question focuses on how abstractions impact privacy configurations. The studies conducted for this particular topic are probably the most closely related ones to past work regarding LSA privacy controls. However, while prior work has shown that users *prefer* having a disclosure abstraction option [Consolvo, Smith et al., 2005; Patil and Lai, 2005], there has been no empirical evidence to suggest that more users will *actually* choose these abstractions when configuring their privacy settings in a real-world location sharing application. After all, choosing to disclose location abstractions is still opting to share some amount of location information (albeit less than the fully descriptive option of sharing one's geographical coordinates). It is entirely possible that, in practice, users will lean towards more conservative disclosure decisions and still choose to disclose nothing. To determine if this is the case, we use a field deployment and a simulated deployment situation to examine end-users' privacy preferences to see how they will actually react to the addition of location abstractions in LSAs.

Our third research topic will explore how to visually present location abstractions in LSAs. Specifically, our intention is to better understand whether different location visualizations can impact a users' *perceived* privacy concerns. In terms of understanding how to actually implement location abstractions in a LSA, past work has not yet addressed these two issues. Consider that nearly all commercial location-aware social applications use maps to visualize users' locations. Since GPS technology implicitly represents locations as geographical coordinates, it is understandable that often the easiest implementation for LSAs is to leverage the numerical properties of coordinates and visually represent locations with pushpins at precise geographical coordinates on a map, as in Figure 5. However, such precise depictions of one's locations can lessen a user's comfort levels about location sharing. We intend to better understand the relationship between location visualizations and perceived privacy concerns.

As part of our exploration of location visualizations, we will look at different ways to include location abstractions. In particular, this dissertation aims to provide a first step towards in understanding the design space for possible visual representations for location information and to understand how different combinations of visual elements influences end-user privacy concerns. Current visualizations of location abstractions are quite limited. When using Google Latitude, users who opt to only disclose city-level information (as opposed to fully precise geographical coordinates shown in Figure 5) will still appear on other users' map, though with very subtle differences. Google states that the distinction is that the marker will be "without an arrow underneath it and [always] without an accuracy circle in the map view. [The] photo icon will also appear in the middle of the city location" [Google, 2009]. As part of the third research question in this dissertation, we will look at a broader selection of visualization candidates and, in particular, compare these visualizations to the current default map-based visualizations to see which is more useful and usable for social location sharing.

Figure 5. Map-based visualization provided for users who opt to disclose a fully precise description of their location using Google Latitude [2009].

The last research question will focus on the potential outcomes one might expect from incorporating abstractions into LSAs. Up until now, we have look at users' reasoning about location sharing, as well as usability issues relating to privacy configurations and location visualizations. Results from all three of these research questions are primarily to

help inform future LSA designs. However, these results are mostly based on *subjective* and *perceived* end-user privacy concerns and utility analyses. In our last research topic, we intend to use a more data-driven approach to examine *actual* privacy costs from engaging in social location sharing and *actual* potential for social interaction (based on historical data). Both of these issues have not yet been explored in past work. The recent flurry of development activity for location sharing applications suggests that at least service providers are convinced that there is indeed meaningful social value to engaging in location sharing behaviors. Redesigning current LSAs to shift from sharing fully precise location descriptions to only sharing location abstraction can potentially require a significant amount of re-coding. Our intention is to supplement our design guidelines with quantitative analyses that support the positive correlation between location sharing, location abstractions, and social interactions. With this information, service providers will be more motivated to consider redesigning their location sharing applications to include disclosure abstractions. However, it is also important to consider the privacy implications of users' decisions about sharing location abstractions. We are specifically interested in seeing how privacy-preserving users' decisions *actually* are, not how privacy-preserving they *intend* their decisions to be. By definition, abstractions are inherently less descriptive, so this would suggest that there is a benefit to incorporating them into LSAs. However, in our work, we intend to better quantify how much privacy location abstractions can really provide for users.

In summary, these four privacy issues form the basis of this dissertation work. Our basic intention is to provide a better understanding of the usability implications for incorporating disclosure abstractions in social location sharing applications. We are interested in probing more than just how these abstractions affect privacy controls. In particular, we will explore the following research questions:

- Q1: How do users *reason* about social location sharing and, in particular, are abstractions a useful construct in a social sharing context?
- Q2: How do abstractions affect users' *privacy configurations* for when they specify their location sharing preferences?

- Q3: How do *visual representations* of location abstractions impact users' sharing preferences and their perception of social utility?

- Q4: What types of *outcomes* can be expected when using abstractions for social location sharing, both in terms of privacy and social utility?

These research questions are explored in six different user studies, which are mapped according to Figure 6.

reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 6. An overview of the four research topics that will be covered in this dissertation, also showing how our six user studies are spread across these research areas.

The first study we conducted examines the properties of social (vs. non-social) location sharing and examines the role of location abstractions in users' sharing preferences. The second and third user study examines how abstractions impact users' privacy preferences for location sharing. The fourth study compares three types of visual representations for location abstractions: text-based, map-based, and time-based. Our final two studies examines both the privacy and utility of location abstractions using a purely data-driven approach. Based on these studies, we are able to synthesize a list of design suggestions for future LSAs on how they can successfully incorporate location abstractions that can address privacy concerns, while optimizing for social utility as well.

We have also designed our studies to explore different styles of social location sharing (see Figure 7). In particular, we look at LSAs that share *current* and *past* locations, as

well as applications that share location information *asynchronously* and *synchronously*. In the first user study, we explore LSAs that support asynchronously sharing of current location. The second user study explores sharing of current location as well, but in a synchronous fashion. The third user study looks at asynchronous sharing of past locations, while the fourth user study examines synchronous sharing of location history. The data for the fifth user study is borrowed from the first user study, so it too studies asynchronous sharing of current location information. The six user study is based on location sharing habits in Facebook status messages, which is also equivalent to asynchronous sharing of current locations.
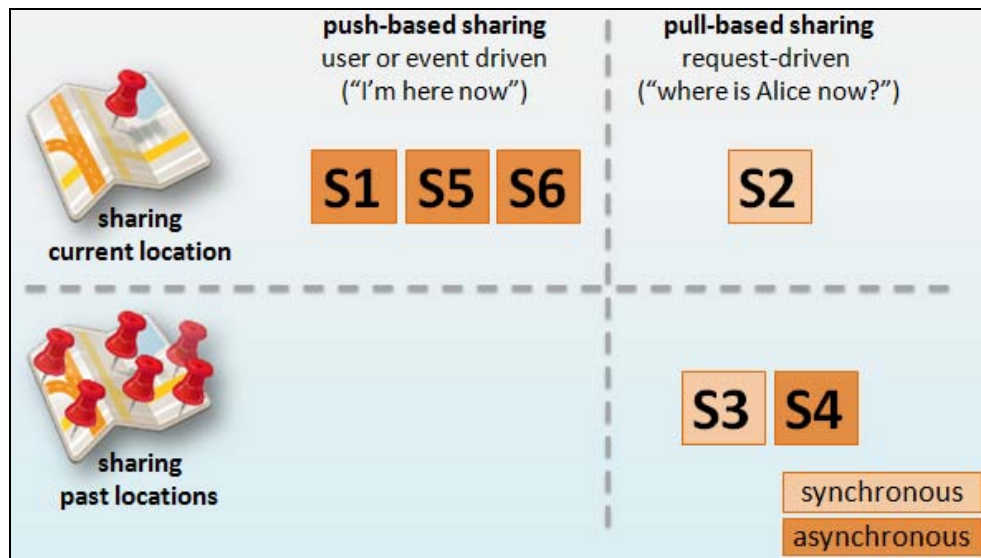
Figure 7. An overview of our six user studies and how they differ in the types of location sharing that they support: sharing current vs. past locations and sharing asynchronously vs. synchronously. Our intention is to explore a diverse set of LSAs in order to obtain a better understanding of how abstractions can impact social location sharing in different contexts.

3.3   Summary

In this chapter, we gave a specific description of the scope for this dissertation. We provided examples for the two types of location abstractions (geographic and semantic)

and introduced the research questions we will be exploring for the rest of this thesis. We also presented a chart that provides an overview of our research agenda, along with how our five user studies fit into those research questions. We briefly presented our four research questions, which is to learn about: 1) how users reason about location sharing, 2) how users configure their privacy settings, 3) how users prefer to visually represent their location information, and 4) what outcomes can be expected when participating in location sharing.

4. Making Decisions About Location Sharing[*]

We begin by examining how users make decisions about location sharing (see Figure 8). Our goal is to better understand if and how location abstractions affect users' reasoning about what types of location information they are willing to socially share with others. Past work has suggested that people often refer to location abstractions conversationally, though this may simply be a result of linguistic constraints, rather than from any explicit user preference. Because location sharing applications often rely on explicit user decisions (e.g., users' privacy configurations), we wanted to more deeply examine how users go about choosing what location information to share and, in particular, whether users reference any types of location abstractions during their decision making process.

reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 8. The four research questions covered in this dissertation. This chapter focuses on the first research question and user study, which looks at how users reason about their location sharing. The goal of this particular study is to understand how location abstractions factor into users' decisions about social location sharing.

---

[*] Portions of the work presented in this chapter was published in [K.P. Tang, Lin, Hong, Siewiorek, and Sadeh, 2010].

4.1   Social-Driven vs. Purpose-Driven Location Sharing

Before the advent of location sharing applications (LSAs), people often obtained location awareness through direct communication channels like phone calls [Weilenmann, 2003], SMS [Grinter and Eldridge, 2001], or instant messaging [Nardi, Whittaker, and Bradner, 2000]. In all of these scenarios, location requests are typically sent from one person to another. With LSAs, we now see a shift in location sharing from previous approaches of one-to-one sharing to current approaches of sharing with many people at once. The push for more information sharing is largely driven by popular micro-blogging and social media sites like Twitter and Facebook, whose users share 50-60 million status updates daily [O'Neill, 2010]. Past literature has shown that these micro-blogging sites are successful in part because they help users build up social capital within their network.

We believe that past instances of LSAs have under-valued this "social" factor. Consider, for example, systems like Reno [Iachello, Smith et al., 2005], WatchMe [Marmasse, Schmandt, and Spectre, 2004], and the Whereabouts Clock [Brown, Taylor et al., 2007]; these LSAs are all motivated by scenarios that emphasize a more utilitarian perspective of location sharing and focuses on activities like coordination and planning. These *purpose-driven* LSAs are in distinct contrast from those that support location sharing within social networks like Foursquare [2009], Loopt [2005], BrightKite [2007], and Locaccino [Sadeh, Hong et al., 2008]. These latter LSAs have motivating scenarios that emphasize the *social* aspects of location sharing, where users might announce their arrival at a location not because others *need* to know, but because it is simply interesting or fun for them to do so. This highlights the fundamental difference between location sharing that is *purpose-driven* vs. *social-driven.*

Past research has primarily focused on what we consider to be purpose-driven location sharing. In this dissertation, we are focused instead on social-driven location sharing and,

in particular, whether users prescribe similar sharing preferences and behaviors between the two types of sharing.

Generally speaking, sharing information within a large social network introduces several interesting properties in which we believe disclosure abstractions could be particularly helpful. We conducted a two-week study collecting actual location traces from nine participants. We focused on two particular aspects of social-driven location sharing. First, we looked at if users chose to share different types of location information, when given different motivations for sharing. Second, we interviewed participants to learn about their privacy concerns for social-driven location sharing and what strategies they used to cope with these concerns. Results from our initial exploration into these issues revealed significant differences between social-driven and purpose-driven sharing. In particular, we found that social-driven location sharing favored semantic location names, blurring of location information, and using location information to attract attention and boost self-presentation.

4.2   Categorizing Existing Location Sharing Applications

In Chapter 0, we discussed several different ways of classifying LSAs, such as by how information is updated (asynchronously or synchronously), who receives the location information (scale), and how information is delivered (push- or pull-based delivery).

We believe that the biggest difference between social-driven and purpose-driven sharing resides in the scale, or the number of people who consume a user's location information. Figure 9 provides a sample classification of some of the more popular commercial and research LSAs. These are arranged along a spectrum, starting with LSAs that primarily support sharing locations with one other person (one-to-one) or with a small group (one-to-few), on up to LSAs that share locations with a large group (one-to-many) or with everyone (one-to-all).
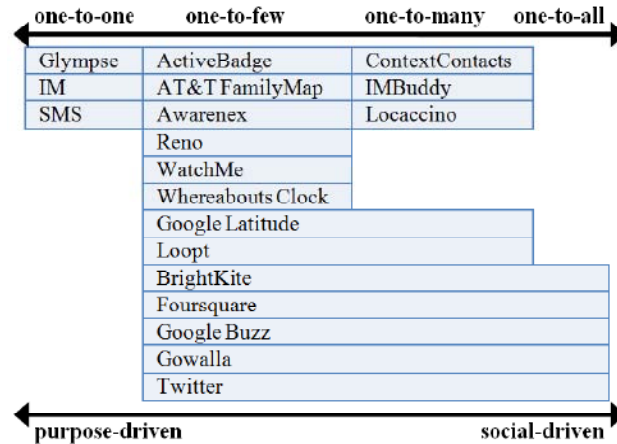
Figure 9. Two ways of describing location sharing apps (LSAs). One is organized by recipient group size. The other is organized by discloser's motivation being purpose- or social-driven.

With one-to-one location sharing, a user's location is shared with one other person. For example, Glympse [2008] lets users send a URL containing their current location to another person. After a specific time period, the map no longer updates. While nothing prevents a user from publicly posting this URL and making it accessible to the world, the original Glympse scenario was to share a time-limited lease of a user's location to one other person.

Other LSAs share users' locations with small (typically homogeneous) groups, like co-workers [Patil and Lai, 2005; J.C. Tang, Yankelovich et al., 2001], family members [Brown, Taylor et al., 2007; Iachello, Smith et al., 2005], or close friends [Barkhuus, Brown et al., 2008; Iachello, Smith et al., 2005; Marmasse, Schmandt, and Spectre, 2004].

There are also LSAs that share location with larger, more diverse groups. These one-to-many LSAs are often integrated with services that provide a relatively extensive social network, like Facebook (e.g., Locaccino [Sadeh, Hong et al., 2008]), instant messaging (e.g., IMBuddy [Hsieh, Tang, Low, and Hong, 2007]), or one's address book (e.g.,

ContextContacts [Oulasvirta, Raento, and Tiitta, 2005]). We also see some one-to-many LSAs opting to use their own application-specific social networks, like Loopt [2005], Foursquare [2009], BrightKite [2007], Gowalla [2009], Google Latitude [2009], and Twitter, with its recently released geo-location feature [Twitter, 2009b].

There are also LSAs that publicly broadcast users' locations so that it is viewable by anyone. In fact, several one-to-many LSAs allow users to publicly share their locations, like Foursquare [2009] and BrightKite [2007]. Alternatively, these LSAs can also be scaled down to function as a one-to-few or even a one-to-one LSA, assuming users proactively adjust their privacy settings so that their location is only shared with specific individuals. It should also be noted that, in practice, users of one-to-many LSAs often have a relatively small social network (like Loopt [2005]), making them more representative of one-to-few location sharing.

The range of one-to-one to one-to-all sharing is important to our framing of purpose-driven and social-driven location sharing. Often LSAs that support one-to-one and one-to-few sharing are purpose-driven sharing, while one-to-many and one-to-all sharing is more social-driven (Figure 9). Thus, to compare these two kinds of location sharing in our study, we use a one-to-one LSA to represent purpose-driven sharing and a one-to-many LSA for social-driven sharing.

4.3   Lack of User Studies for Social-Driven Location Sharing

Our expectation is that social-driven LSAs elicit significantly different privacy concerns than purpose-driven LSAs. Lederer et al [2003] and Consolvo et al. [2005] explored related issues in their work. Using ESM and hypothetical location requests, they found that the primary factor for location sharing was based on who sent the request. Why the request was sent also factored into users' decisions about what information to share, albeit to a lesser degree [Consolvo, Smith et al., 2005]. For our purposes, we consider Consolvo and Lederer's work as primarily focused on one-to-one (purpose-driven) sharing, where users share their location to only one other person.

We believe the type of sharing described by Consolvo and Lederer is markedly different from one-to-many (social-driven) sharing. Barkhuus et al.'s Connecto [2008] comes a bit closer to this type of sharing, but still focuses on what we consider small-group (one-to-few) location sharing between close friends. Large-group (one-to-many), social-driven location sharing scenarios introduces more privacy concerns than small-group sharing because there are inherently more relationship types to handle. In Facebook, prior work has shown that users' social networks mostly consist of "loose" social connections or acquaintances [Donath and boyd, 2004; Wellman, Haase, Witte, and Hampton, 2001]. We expect that location sharing within these groups will have vastly different privacy implications than when sharing locations with just close friends or with one other person.

4.4   Why One-to-Many Sharing Introduces Privacy Concerns

In one-to-one location sharing, the user's decision is simple: is the user comfortable telling this specific person her location? For one-to-many sharing, the decision is more complex: what may have been okay sharing with one person may not be okay sharing with many people. There are three reasons why large-group sharing might differ: (1) there is a larger variance in who receives the information, (2) there is a different motivation for sharing, and (3) there is a different expectation of plausible deniability.

4.4.1   Who is the Location Information Being Shared With

Large-group sharing involves disclosing location information to a diverse social network. Currently, large-group LSAs are integrated with an online social network like Facebook. The size of these social networks is often several orders of magnitude larger than offline networks [Gross and Acquisti, 2005]. Online social networks often also include several weak social ties [Donath and boyd, 2004; Wellman, Haase, Witte, and Hampton, 2001] and weaker ties suggests that there will likely be a large variance in how much the user trusts their social network with the user's personal information.

These features have significant privacy implications for location sharing. The success of Facebook is indicative that users are relatively comfortable sharing the same status information with everyone in their online social network (i.e., people of varying tie strength), but it is unclear if the same holds true for location sharing. For example, users may be comfortable telling their close friends that they are "at the movie theater", but are they equally comfortable sharing that with everyone else in their network? Will users employ different strategies for sharing location abstractions when comparing social- vs. purpose-driven location sharing?

4.4.2   Motivations for Location Sharing

For most one-to-one LSAs, the disclosure process begins with the requester. For example, Bob wonders where Alice is, so he sends a request to Alice asking for her location. This request-response model allows users to decide what location information to share using information like: (1) who is receiving the information, (2) what is the most likely reason for why the request was sent, (3) what would be most useful, given this reason, and (4) is the user comfortable sharing that level of location information [Consolvo, Smith et al., 2005].

We consider this type of location sharing as *purpose-driven location sharing* since the requester most likely has a specific need for the user's location. This kind of behavior is used in many scenarios motivating prior LSAs (e.g., Reno [Iachello, Smith et al., 2005], the Whereabouts Clock [Brown, Taylor et al., 2007]) and in past ESM studies [Consolvo, Smith et al., 2005]. In past diary studies, it was shown that 85% of location requests were for pragmatic reasons, including coordinating meetings, arranging transportation, sending reminders, providing roadside assistance, checking for availability, and asking for estimated time of arrival (ETA) [Reilly, Dearman, Ha, Smith, and Inkpen, 2006]. Consequently, in purpose-driven location sharing, the disclosure decision is often a pragmatic one: does the reason warrant a disclosure and what would be the most useful location information for this purpose?

On the other hand, large-group location sharing is better framed as *social-driven location sharing*. Current disclosure behaviors on social networks sites like Facebook reveal that users generously share their information [Gross and Acquisti, 2005]. Prior work has shown that this information exchange helps build up social capital [Ellison, Steinfield, and Lampe, 2007]. Similarly, we believe that large-group location sharing can enhance peripheral awareness, which has shown to help promote and sustain social capital within one's network [Resnick, 2001]. In other words, we expect that, just as general-purpose information sharing is driven by social capital, large-group location sharing will also be driven by similar motivations like social capital.

Generally speaking, our observations of past LSAs reveal that purpose-driven location sharing is often aligned with one-to-one and one-to-few location sharing. Social-driven location sharing, on the other hand, is closely aligned with one-to-many location sharing. It is important to note that the distinctions between purpose-driven and social-driven location sharing can be somewhat fluid. For example, consider a mother who is wondering if her son has arrived at his spring break destination. Her request (and her son's subsequent location disclosure) would fall under purpose-driven location sharing. However, it is possible that there is some hint of social capital involved since the mother may now feel more in-tuned with her son's activities (i.e., it contributes to her peripheral awareness). Despite this effect, we would argue that the son's primary motivation for sharing his location is most likely purpose-driven, as her son probably reasoned that his mother *needed* to know the information (e.g., for okayness checking [Iachello, Smith, Consolvo, Chen, and Abowd, 2005]), as opposed to primarily asking just for the sake of curiosity.

Continuing this example, consider if the son had shared his location with his online social network. In this case, no individual person is requesting his information, but he still chooses to share it. We would argue that, in this case, his decision to share his location is mostly to increase his social capital and, as a result, his social network is more aware of his activities as revealed through his location information.

4.4.3 Expectations of Plausible Deniability

Prior work has suggested that LSAs should support plausible deniability so that users can "stretch the truth" [Iachello, Smith et al., 2005]. However, in field studies of LSAs that use one-to-one (purpose-driven) sharing, actual occurrences of outright deception are relatively uncommon, though use of location blurring does sometimes occur [Consolvo, Smith et al., 2005; Iachello, Smith et al., 2005].

For one-to-many (social-driven) location sharing, we expect that there may be more incentives to exercise deception. Evidence already exists in online social networks [boyd, 2004]. Social psychology literature also informs us that people often tell self-centered lies to make themselves look or feel better, or to protect themselves from embarrassment or disapproval [DePaulo and Kashy, 1998]. This type of behavior is especially prevalent in casual relationships (e.g., acquaintances), as opposed to close relationships (e.g., family) [DePaulo and Kashy, 1998]. Since one-to-many location sharing most likely involves more casual relationships, users may end up choosing to exercise plausible deniability when sharing their location.

4.5 Research Questions

By conducting a comparative study, we can contrast users' privacy concerns for social- vs. purpose-driven sharing. In particular, we will focus on two research questions:

- Does social-driven location sharing result in different location sharing decisions?
- What privacy strategies are used in social-driven (vs. purpose-driven) location sharing scenarios?

4.6   User Study

To address these research questions, we conducted a two-week user study in November 2009 with ten participants, all of whom were recruited through a university-wide mailing list. One participant dropped out midway due to scheduling conflicts. Participants ranged from 18-46 years old (μ=27.1, σ=8.3); three were female. Two-thirds were either undergraduate or graduate students; the remaining participants were university staff members. Participants were evenly split between those affiliated with technical (e.g., natural sciences, engineering) and non-technical fields.

4.6.1   Part 1: Entrance Survey

Participants completed a 10-min online survey to collect basic demographic and social network information. We intentionally did not ask include privacy to avoid biasing participants later. For their social networks, participants provided examples (names) for four relationships: family members, acquaintances, managers/bosses, and close friends. We told participants that their examples must live in the same city. This way we control for geographical distance and avoid having that influence participants' location sharing decisions. The names that were collected were used when creating scenarios for later on in the study.

4.6.2   Part 2: Location Data Collection

Participants were given mobile phones (Nokia N95s) to carry for two weeks and were required to use the N95s as their primary mobile phone. This helped to incentivize them to keep the phone sufficiently charged at all times.

The phones were equipped with location-logging software that was written in C++ and was used to collect participants' actual location traces (the same software used in [Benisch, Kelley et al., 2008]). The software ran continuously in the background (without

user input), using both GPS and Wi-Fi positioning technology. To reduce power consumption, the application used the phone's accelerometer to selectively sample location information only when significant movement was detected. The N95 phones have built-in 3D accelerometers that can sense acceleration along three dimensions at a rate of about 40 samples per second. The software records a moving average (using a window of 2 minutes) of the phone's acceleration along each axis. If the phone passed a threshold of 0.1 g's (after accounting for gravity) within that window, then the phone would begin recording GPS readings every 15 seconds. To accurately record indoor locations, when GPS readings tend to be imprecise or non-existent, the application also tracked nearby WiFi access points by recording their MAC addresses and their corresponding signal strengths every 3 minutes. All the recorded information was stored locally on the device by appending to the file whenever the application was actively tracking the user's location. Readings were always recorded with its associated timestamp in order to sequentially sort of the order of the sensor data. We provided daily email reminders for participants to upload their location data each day of the study, which they did by plugging in the device to their computer with a USB cable that we provided. These files were uploaded to a web application that we created so that these files would feed directly into our backend databases.

It is worth noting that the settings for our location tracking software were determined empirically, after conduct several small scale experiments measuring the device's battery life. We chose a smaller sampling frequency (i.e., sampling less often) for WiFi readings because accessing the device's WiFi sensor consumed considerably more energy than the GPS sensor. There is one exception to this rule, which is when the device is initially trying to acquire a GPS lock on its position. However, after this initial task, subsequent readings are relatively inexpensive, in terms of energy consumptions. Thus, by using motion-triggered location tracking, we hope to minimize the amount of energy used to acquire GPS locks.

When significant motion was detected, the GPS unit began recording every 15 seconds until the GPS signal disappears. The application recorded Wi-Fi MAC addresses every 3

minutes if the GPS signal was too weak. All location traces were stored locally on the device. We provided daily email reminders for participants to upload their location data each day of the study.

We acknowledge that there are some shortcomings to our automated data collection. But, by doing so, we had a continuous record of participants' location data, with little to no additional effort from participants. This is especially helpful for places where the participant stops by for only a short time. Manual data collection (e.g., like with ESM) would require interrupting the user and potentially risking large gaps in the location trace if users ignored the prompts.

4.6.3 Part 3: Location Sharing Interviews

Before each interview, we analyzed each participant's location trace. We used Skyhook's API [Skyhook Wireless, 2003] to translate WiFi readings into GPS coordinates. We then computed the distance and speed between adjacent coordinates to determine if the participant was moving. Places that the participant stayed for more than five minutes were marked as "significant". During the hour-long interview, participants completed the following three steps for each location marked as a significant place (Figure 10):

- Describe the place, using up to eight labels
- Given a hypothetical purpose-driven location sharing scenario, choose what label to share and explain why
- Given a hypothetical social-driven location sharing scenario, choose what label to share and explain why

We chose to use hypothetical sharing scenarios instead of actual location disclosures to other people. This decision was primarily to protect participants from unintentionally sharing sensitive locations. To help ground the scenarios for our participants, each scenario referred to a specific person using names obtained at the start of the study. We also asked participants to think of up to eight labels upfront to help ensure that they

carefully considered which location name to share. Interview responses also suggest that participants were thoughtful in their decisions.

Figure 10. An example webpage used in the study. (top) Map reminds participant of a place they visited. They first write labels to describe the place. Next, we show two hypothetical sharing scenarios, randomly ordered. (middle) In purpose-driven scenarios, they read a randomly generated scenario, choose label(s) to share, & describe recipient's familiarity with the place. (bottom) In social-driven scenarios, they see how locations might appear in a social network site & pick label(s) to share.

For each significant place (as described by a timestamp & map, Figure 10a), participants responded to both purpose-driven and social-driven scenarios (randomly ordered). For each location sharing decision, participants were asked to explain to the interviewer their rationale. To mimic purpose-driven location sharing, we had eight hypothetical scenarios in which the request for the participant's current location was motivated by a specific reason. For example, one scenario was: "While you're at this place, Maria (your

roommate) contacts you. She has lost her keys and would like to meet you now to borrow your keys to the apartment now" (Figure 10b). Each scenario refers to a specific person (Maria) and relationship type (roommate), which reflects the one-to-one aspect of purpose-driven sharing. These scenarios are randomly generated by changing the location requester's identity. If a scenario does not make sense (e.g., a manager is looking for your apartment keys), then another scenario is randomly generated. For social-driven location sharing, we presented participants with a screenshot showing how their location might appear on a social network site (Figure 10c).

At the end of the study, participants completed a survey that measures privacy concerns and use of social network sites. Participants were then compensated with a $30 gift card.

4.7 Results

In total, we identified 98 unique significant places from 29,490 recorded location readings from the N95 phones. Each participant visited $\mu=10.9$ unique places ($\sigma=2.2$).

4.7.1 Place Labels

Using a bottom-up approach, we classified all the labels that participants chose to share under both the purpose-driven and social-driven sharing scenarios. Earlier work classified labels as relating to a place ("home") or an activity ("shopping") [Iachello, Smith et al., 2005]. Others have looked at labels as a geographical hierarchy, ranging from street address ("123 Main St.") to neighborhood ("Brooklyn") to city & state [Consolvo, Smith et al., 2005]. Barkhuus's work used four categories: geographic, place-based, activity-based, or a mix of these three [2008].

We felt that these categories were too broadly defined for our purpose. Using similar categories in Lin et al.'s work [2010], we settled on a more detailed taxonomy (Table 3).

In particular, we used a more complete classification scheme for semantic place names that includes personal names ("my home"), functional names ("restaurant"), activities ("shopping"), and public businesses ("Starbucks"). Categories labeled as "specific" vs. "non-specific" refer to when a place name is more precise (e.g., there are several "restaurants", but fewer "Indian restaurants") or is unique (e.g., there is only one "my home", but there are more than one "friend's home").

We also extended the geographical category to include room, floor, and building. This change is mainly since our participants often visited a local university campus, which includes this level of granularity. Note that place labels can fall under multiple categories, so total percentages may exceed 100%. For example, "restaurant@5th & 2nd" counts as both "semantic, functional, non-specific" (restaurant) and as "geographic, street/intersection" (5th & 2nd).

Across 98 unique significant places, participants provided 505 place labels, (μ=5.15 labels/place, σ=1.57). 57.03% of the labels were geographic; 42.97% were semantic. Overall, participants shared more semantic names than geographic names. For purpose-driven sharing, 69.39% of the labels were semantic names vs. 40.20% geographic names. For social-driven sharing, 77.55% were semantic names vs. 25.71% geographic names. Social-driven sharing used significantly more semantic names than in purpose-driven sharing (χ^2=27.74, p<0.001). Considering only semantic names, social-driven sharing also had a significantly different distribution (χ^2=23.68, p<0.005): social-driven sharing favored labels with activity and personal names over functional and public business names.

4.7.2   Location Sharing Decisions

Prior work has found that users will choose to share their location at whatever level of detail is most useful, or to share nothing at all if the request is inappropriate [Consolvo, Smith et al., 2005]. Given that our scenarios are purpose-driven, we were interested in

Type of Place Label	Examples	Purpose-Driven Location Sharing (%)	Social-driven Location Sharing (%)
Semantic	---	**69.39**	**77.55**
Personal	---	**12.24**	**17.35**
Non-specific	friend's house	2.04	4.08
Specific	my home, my office	10.20	13.27
Functional	---	**17.34**	**14.28**
Non-specific	restaurant, library	10.20	9.18
Specific	Indian restaurant	7.14	5.10
Activity	---	**16.32**	**31.35**
Activity only	in class, shopping	7.14	19.39
Activity@location	shopping @ Walmart	6.12	7.14
In transit	on my way home	3.06	4.82
Public business	---	**23.47**	**15.30**
Not unique within city	Starbucks, Barnes & Noble	10.20	5.10
Unique within city	Lewis Salon	13.27	10.20
Geographic	---	**40.20**	**35.71**
Room	<building name> <room number>	5.10	0.00
Floor	<floor number> <building name>	4.08	0.00
Building	<building name>	23.47	15.31
Address	500 Main St	6.12	0.00
Street/Intersection	Main St & 1st Ave	11.22	4.08
Neighborhood/Region	Downtown	6.12	5.10
City	San Jose, New York	4.08	11.22

Table 3. Taxonomy for place labels that includes both semantic and geographic place names. Breakdown of labels for each of the 98 unique places obtained from our participants over a two-week period for both purpose-driven and social-driven location sharing. Note, total percentages exceed 100% since place labels can be classified under more than one category.

whether participants would unilaterally provide the most precise location label (typically a geographical name), or if they still opt to selectively share their location information. To investigate this issue, for each purpose-driven sharing scenario, participants provided a familiarity score (5-point Likert scale; 1=completely unfamiliar) to describe how familiar the requester was with the participant's shared location. When recipients were unfamiliar with the location (scores<3), participants opted to share more hybrid labels (using both geographic and semantic labels). With higher familiarity scores ($\geq$3), participants opted to share labels that contained only semantic place names. This difference was statistically significant (G^2=13.32, p<0.002) and indicated that our participants selectively decided what to share based on the recipient's familiarity with the place.

In the social-driven sharing scenarios, we looked at if certain types of locations led to participants' preferring geographical or semantic labels. We found that, for locations identified as home & work, participants unilaterally used semantic place names. When participants were at public locations (e.g., grocery store, Starbucks), they were more likely to share semantic place names, followed by functional place names and public business names. We also examined whether certain properties of locations led to specific sharing behaviors. We found that when participants visited public locations typically having a lot of people (e.g., grocery store vs. salon), they preferred sharing functional place names.

4.8   Discussion

Our main research goal is to compare purpose-driven and social-driven location sharing. Information sharing has generally shifted from being one-to-one to now being one-to-many. In addition, information sharing is often tightly integrated with large social networks that span several relationship types. The diversity and size of these networks lead to several potential privacy concerns, particularly when it comes to sharing sensitive information like one's location information. By comparing purpose-driven and social-driven location sharing, we hope to better understand users' privacy concerns and preferences through their decisions about what locations they share under each condition.

4.8.1   Differences in Location Sharing Decisions

We found that participants share different place names for social-driven location sharing. When considering only three types of labels (geographic-only, semantic-only, and hybrid – a mix of geographic and semantic names), we found that social-driven sharing led to more semantic-only place names (39.80% vs. 64.29%, $p<0.01$) and fewer hybrid place names (29.59% vs. 13.27%, $p<0.005$). Generally speaking, hybrid names are more descriptive since they provide both geographic and semantic information. Sharing fewer hybrid names suggest participants prefer the ambiguity of semantic place names. There

was no difference for geographic-only names (30.61%, purpose-driven vs. 22.45%, social-driven).

Our distribution of geographic, semantic, and hybrid names is similar to the distribution found in [Barkhuus, Brown et al., 2008]. However, in our study, we can also examine labels that participants did not choose to share. In 64.29% of these cases, participants shared semantic place names (for social-driven sharing) and explicitly did not pick a geographic name that was listed in their list of possible place labels. This finding suggests that participants do make deliberate decisions when choosing to share a particular type of label over another.

When asked why they made their selections, participants cited two main factors: privacy concerns and attracting attention. For example, P5 reported choosing a label as a way to advertise to others that he might be nearby to them: "If any of my friends happen to be around me, then I will probably meet with them." This is similar to Weilenmann's observation that place is sometimes used to express availability [2003]. In her study, she examined one-to-one (purpose-driven) location sharing. In our study, we confirmed a similar use of location information for one-to-many (social-driven) location sharing as well.

We also observed that social-driven location sharing decisions were influenced by impression management. For example, P3 reported that "being at Mad Mex [a local restaurant] is pretty cool and I want people to know that." This finding suggests that, for social-driven location sharing, participants use location information as an indirect way to enhance their self-presentation so that they appear more interesting to others in their social network.

4.8.2 Perceived Privacy Strategies

Based on their Westin scores [Kumaraguru and Cranor, 2005] obtained at the end of the study, most participants were privacy pragmatists (5/9), one was privacy unconcerned, and two were privacy fundamentalists. This classification suggests that most participants had balanced privacy attitudes and would be willing to forego some privacy if there is a clear benefit. Since our study used only hypothetical scenarios, one might expect our participants to exercise highly conservative location sharing behaviors. Instead, we observed only slight use of location blurring using three different strategies.

The most often used method was to leverage "insider knowledge" to obscure one's actual location. This strategy provides users with plausible deniability for providing less precise location information. For example, P6 shared that he was "at Giant Eagle" (a local grocery store chain) and said that he chose to share this because "for people who know where I live, they can figure out which Giant Eagle I am at, otherwise, they won't know".

Similarly, P5 shared that he was "at INI" (a university building) because "if I say INI, classmates will know where I am, but, for other people, they will have no idea what INI is." This suggests that participants are actively deciding to blur their location.

It is important to note that the location blurring we observed is a relatively minor type of deception. When deciding what to share, participants were not precluded from lying and they could have opted to share fake labels. However, during our interviews, none of the participants chose to share outright false location information. Participants could have also hidden their true location by blurring at the city or state level. However, for social-driven sharing, we found no evidence of blurring at the state level and only 10.2% of all place labels used blurring at the city level (20% of these occurred when one participant was traveling out-of-state).

Our supposition is that participants' preferences for relatively minor location blurring are related to our previous observation that location sharing is often used for impression

management. By opting to share a place name that is somewhat precise ("Giant Eagle"), as opposed to one that is fully precise ("Giant Eagle @ Center Ave"), participants can still appear as though there are actively involved in contributing to their social network's overall social capital. If they opt to share an overtly vague place label (e.g., "Pennsylvania"), then it may come across as though they are intentionally being socially reclusive.

A second privacy strategy that we observed was where participants hid their location information by opting to share their current activity instead of their current location. In fact, many participants cited that they were generally more comfortable sharing activity information: "I feel like sharing activity should not be a problem" (P4), "I'd rather say what I am doing than that I'm at a certain place" (P2), and "In general, I don't mind telling others what I'm doing" (P7). This is different from prior work which has stated that users opt to share activity in order to be *more* descriptive about their current state [Iachello, Smith et al., 2005; Weilenmann, 2003]. Our findings suggest instead that participants are opting to share activity information for plausible deniability reasons. In other words, sharing one's activity is perceived as *less* descriptive than sharing one's location.

Of all the activity-related semantic names (31.35%), six common types of activities accounted for 78.26%: in class, working, with family, eating, in meeting, and shopping. Other activities were also shared (e.g., "getting a haircut, "dance practice"), but these were used by specific participants. Further work is needed to determine if these common categories can be generalized for other users.

The third privacy strategy that we observed was that participants all seemed to highly value their friends' location privacy. For example, while P5 was at her friend's apartment, she explained that "I'm uncomfortable sharing with people where I am at, since it's someone else's place." P8 had similar concerns: "Sharing a friend's name [in my location] is too much. People don't need to know her name." These responses suggest that participants are highly conscientious about sharing their friends' location. There are

two possible motivations for this privacy behavior: (1) sharing a friend's name reveals the participant has a relationship with that person, or (2) sharing the friend's name reveals not only the participant's location information but their friend's as well. This finding is interesting given that prior work has found that social network users are often quite causal about sharing their friends list [Donath and boyd, 2004]. By attaching location to a friend's identity, our participants seem to have adopted a more conservative perspective.

These three privacy strategies, as observed through participants' interview feedback, were much more prevalent in social-driven location sharing scenarios. It should be noted though that purpose-driven sharing also practiced these blurring techniques to some degree. However, the critical difference is in the motivation behind using these strategies. In social-driven sharing, participant reported using privacy strategies in order to "hide" or blur their true location. In purpose-driven location sharing, participants blurred their true location primarily to convey their unavailability: "My manager doesn't need to know where exactly I am, so I will just tell him I'm at a restaurant [as opposed to the name of the restaurant]." (P6).

4.9   Implications for Future Location Sharing Applications

Our study is also only an initial exploration into the differences between purpose-driven and social-driven location sharing. We designed our study to compare two extremes of the spectrum: one-to-one purpose-driven location sharing and one-to-many social-drive location sharing. There are certainly other possible combinations worth exploring in future work. For example, crises like the U.S.'s Hurricane Katrina demonstrate the need to have one-to-many purpose-driven location sharing, where people can broadcast their location as an indication to their social network that they have reached a safe location.

Despite this limitation, our findings show that there are significant differences between purpose-driven and social-driven location sharing. These results have several design implications for future LSAs. First, LSAs should consider which type of location sharing

they are primarily supporting. Purpose-driven sharing resulted in users sharing different types of location information, compared to social-driven sharing. These differences have clear implications in terms of what data types to support and what type of visualizations to have. For example, social-driven location sharing showed a preference for sharing activity, not just location information. Semantic names were also generally preferred for both purpose-driven and social-driven location sharing. In addition, locations shared in social-driven scenarios were significantly less suited to map-based lookups than purpose-driven scenarios ($p<0.0001$). This result suggests that LSAs might consider other location displays instead of pushpins on a map, which the default visualization used in LSAs like Google Latitude [2009] and Locaccino [Sadeh, Hong et al., 2008].

Another important finding from our data is the factors involved in users' location sharing decisions. In social-driven location sharing, the identity of the requester is ambiguous, making a utility-based decision process (like that suggested in prior work) impractical. Instead, we found that, for social-driven sharing, users attempted to balance between maximizing their social capital while protecting their own privacy. In particular, users want to share information that is interesting, enhances their self-presentation (impression management), and/or leads to serendipitous interactions. Social-driven LSAs can leverage this information by playing to these factors in order to encourage users to share their location. This will, in turn, enhance peripheral awareness within users' networks and allow them to reap the social benefits of location sharing.

4.10 Summary

In this chapter, we described social-driven sharing, distinguishing it from past examples of what we refer to as purpose-driven location sharing. We also explored the differences between these two types of sharing by conducting a two-week comparative study with nine participants.

We found significant differences in terms of users' decisions about what location information to share, their decision process, and their intentions behind blurring their location information. In particular, we found that social-driven location sharing favors semantic location abstractions over geographic abstractions. The types of semantic abstractions, however, tend to favor activity-based labels (e.g., "shopping", "driving", etc.) and personal labels (e.g., "home", "work"), both of which can be computationally difficult to implement in location sharing applications. For example, to support activity-based labels, an application would need to have some sort of activity-based recognizer which could map a specific address (e.g., "501 West Waterfront Dr, West Homestead, PA") to a specific activity (e.g., "(grocery) shopping"). The translation between space and activity can also be a non-trivial engineering task. To support personal labels is likely to be less difficult, as most people tend to visit only 1-2 places per day [Lin, Xiang, Hong, and Sadeh, 2010] and only a subset of these places are likely to be routine (i.e., regularly visited) places. However, we found that the preference for semantic abstractions was significant (for both social and purpose-driven sharing); this is an important consideration for those developing future location sharing applications. We also found that many of the blurring strategies used by our participants were implicitly supported through the use of either geographic or semantic abstractions as well. These results suggest that users did consider abstractions when making decisions about what location information to share. Moreover, we observed that our participants used abstractions both as a privacy mechanism and as a tool for having more meaningful social interactions with those in their social network.

5. Privacy Configurations, Part 1[*]

For this chapter, we move on to the second research question, which is to examine how location abstractions impact end-user privacy configurations (Figure 8). Most privacy-related studies regarding location sharing has been focused on the design and use of specific types of privacy controls. Much of the empirical work though has been limited to laboratory user studies. In our work, we extend this work by exploring how location abstractions are used *in practice*, through a real-world deployment of a location-aware social application. We also conduct a comparative study to better understand how explicitly excluding options for location abstractions can impact end-user privacy configurations. These studies allow us to better isolate the effects of including abstractions in location-aware social applications.

reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 11. The four research questions covered in this dissertation. This chapter focuses on the second research question and user study #2, which looks at how location abstractions are utilized in privacy configurations. In this particular study, we examine privacy policies for applications that share current location information synchronously.

---

[*] Portions of the work presented in this chapter was published in [Hsieh, Tang, Low, and Hong, 2007].

5.1   Motivation

In the previous chapter, we systematically differentiated two types of location sharing: social-driven location sharing and purpose-driven location sharing. We theorized several reasons why social-driven location sharing should present different types of privacy challenges for users. Upon a side-by-side comparison of sharing preferences and behaviors, we did indeed find several ways that social location sharing differed from non-social location sharing. In particular, we found that, for social location sharing, users are much more cognizant of potential privacy issues and can make careful decisions about what sorts of location information they should share with others.

This finding forms the basis for the rest of the studies in this dissertation. First, we have shown that there is a strong preference for sharing location abstractions in social-driven location sharing scenarios. Second, we have evidence suggesting that social-driven sharing is a good candidate for studying end-user privacy concerns, as users are more likely to approach this task using privacy-related justifications to explain their sharing preferences (when compared to purpose-driven sharing). This finding suggests that social location sharing is a suitable basis for further exploration into end-user privacy concerns and, in particular, to examining privacy issues relating to the rest of the process for social location sharing (Figure 6).

In this study, we focus on understanding how location abstractions impact users' sharing preferences, in terms of their privacy configurations. Including disclosure abstractions will inherently make privacy configuration interfaces more complicated, as there will be more options that users must choose from when deciding how to define their privacy rules. In addition, it is not clear whether abstractions provide enough plausible deniability for users so that they will prefer to share that fidelity of location information, as opposed to simply not sharing any location information at all.

To address this issue, we needed to assess users' privacy concerns in a realistic environment where their location information is actually exchanged with other users and,

hence, their privacy configurations are vitally important to controlling how they want their location to be shared. This led us to create a context-aware instant messaging (IM) system called IMBuddy.

5.2   IMBuddy System Design

IMBuddy was designed to support disclosure requests for several types of context information, including interruptibility, current task (as indicated by the title of the active window on the desktop), and location. However, for the purposes of this dissertation, I will treat IMBuddy as just a location-aware IM system.

There are three parts to IMBuddy: an IMBuddy AOL Instant Messaging Bot (AIM) Bot called "imbuddy411" (implemented using JAIMBot, an open-source, Java-based AIM library [Oster, 2005]), an IMBuddy server, and an IMBuddy client running on a WiFi-enabled device. Any AIM user can request location information for any IMBuddy user. To initiate a request, an AIM user types a text command in a chat window to imbuddy411. For example, he can type "whereis X" to get X's current location, where X is the screenname of the IMBuddy user (Figure 12, step 1). imbuddy411 passes this request to the IMBuddy server, which then communicates with the appropriate IMBuddy client to retrieve the user's location information (Figure 12, step 2). The client relays its location information back to the IMBuddy server, which then send a reply back to the AIM user (Figure 12, step 3) and notifies the IMBuddy user that a location disclosure has occurred (Figure 12, step 4). The level of location information that the IMBuddy discloses is dependent on the IMBuddy user's initial privacy configuration. All location requests and subsequent responses are stored in a MySQL database on the IMBuddy server. This lets the server share the most recent location information in the event that the IMBuddy user is offline when an AIM user sends a location request.
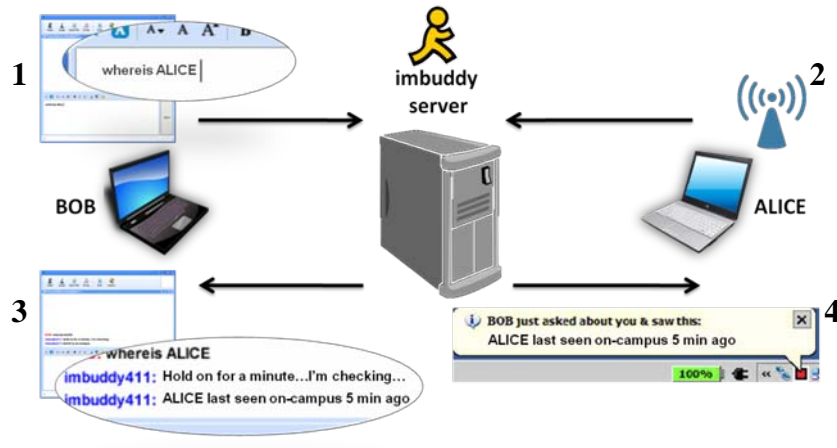
Figure 12. (1) Bob queries for Alice's current location by typing "whereis ALICE" to imbuddy411; (2) imbuddy411 sends the location request to the server, which then finds Alice's IMBuddy client and waits for her laptop to report back its location information; (3) after receiving and filtering Alice's location information (based on her privacy settings), imbuddy411 sends a reply to Bob's request; and (4) Alice's client notifies her that her location information has just been shared with Bob.

The IMBuddy client software runs as a background process that collects location information using a WiFi-based algorithm. Because our participants are college students, the first pass for location positioning checks if users are on or off campus by determining if their IP address is within the university's subnet. For off-campus locations, IMBuddy uses a web service to identify the user's current city based on their IP address. To provide more precise location information, IMBuddy relies on Place Lab [LaMarca, Chawathe et al., 2005] to sense nearby wireless access points. When the client application sees a new set of wireless access points, it prompts users to provide a location tag. Later, IMBuddy will use Place Lab to recognize when the user returns to that location, so that it will not need to prompt the user again. The IMBuddy client also provides notifications whenever location information is shared (e.g., Figure 12, step 4).

It should be noted that IMBuddy supports several other types of feedback mechanisms, in addition to the real-time notification shown in Figure 13a. Examples of these include: 1)

IMBuddy's disclosure log where the system provides an abbreviated summary of how many disclosures has occurred in the past six hours as well as a webpage showing the complete history of information disclosures (Figure 13a-c); 2) IMBuddy's social translucency reminders which alert the discloser about their most recently shared information each time a conversation is started with one of their IM buddies (Figure 12d); and 3) IMBuddy's peripheral notifications which alert the user as to whether they are online (in which case their current location would be shared) or offline (in which case their last known location would be shared) (Figure 12e). These feedback mechanisms served as important privacy-related features to further address potential end-user privacy concerns, beyond just the privacy configuration level. However, the design rationale for including these feedback mechanisms and the analysis of users' preferences for these features is outside of the scope of this dissertation. We refer the reader to [Hsieh, Tang, Low, and Hong, 2007] for more details.
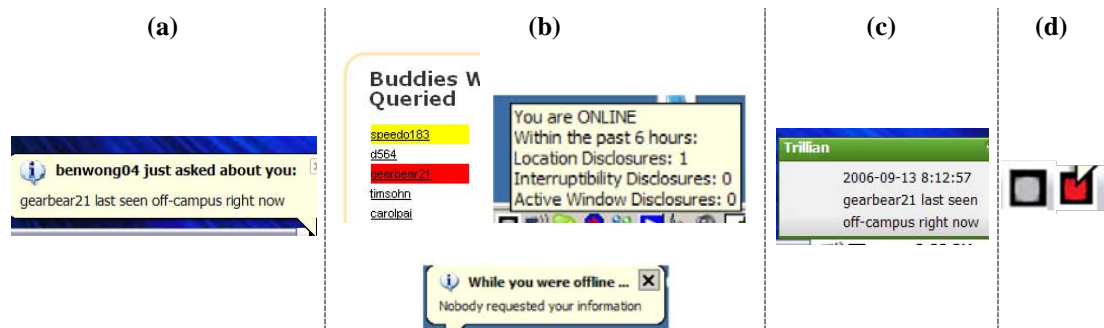
Figure 13. Six examples of the different feedback mechanisms supported by IMBuddy: real-time notifications (a), disclosure history (b), social grounding/translucency (c), and peripheral status notifications (d).

In terms of location disclosures, IMBuddy's privacy controls support three levels for location sharing (Figure 14, left). The lowest disclosure level is "none", which results in imbuddy411 sending a reply of the form "no information is available for X", where X is the IMBuddy user whose location is being request. The highest disclosure level shares the user's self-specified location tags. This is similar to the user-created labels in systems like Reno [Iachello, Smith et al., 2005]. The middle disclosure level shows whether the user is

on or off campus and also provides a city and/or neighborhood description, if the user is off-campus. Thus, IMBuddy's three-tier disclosure scheme included both geographic abstractions (in the middle disclosure level where city and neighborhood information is shared with others) and semantic information (in the top disclosure level where personal labels are shared with others).

5.3   User Study

IMBuddy was deployed for four weeks. It was important that the deployment be significantly longer than just a few days, which is the typical duration of past field studies for location-aware systems. Moreover, a longer deployment period allows us to more thoroughly probe users' privacy concerns about location sharing.

Throughout the study, participants provided self-reports of their privacy comfort level for sharing their location information and the perceived appropriateness of the location information that had been shared. To collect this feedback, participants were interviews three times during the four-week period. The first session occurred immediately before the four-week study began. At this session, we introduced the IMBuddy system to participants, and asked that they set their privacy preferences for disclosing their location information. IMBuddy used a group-based approach adapted from prior work by Patil and Lai [Patil and Lai, 2005] (Figure 14, left). Because IMBuddy is an IM-based system, participants were required to specify disclosure settings for all of the screen names in their buddy list. Initially, each participant's complete buddy list is classified as a 'default' group in their privacy settings, which, by default, discloses the minimum amount of location information (i.e., nothing is disclosed). Participants are asked to modify the default privacy settings as they see fit. They are allowed to create as many or as few disclosure rules as they want and/or feel comfortable with. New groups are created by dragging a new from the 'default' group to a new group. The only requirement given to participants is that they must include one 'default' group in their disclosure configuration. If an unknown AIM user (i.e., a screen name which does not show up in any of the

participant's manually created disclosure rules) requests location information using IMBuddy, then IMBuddy would use the settings in the 'default' group when deciding what level of location information to disclose.

Through formative user tests, it was found that a group-oriented view that lists the group's privacy information in a vertically-oriented container is preferred. Users cited the similarity of this layout to that of existing IM buddylist views as the primary reason for this preference. Within each group's container, drop-down controls let users modify the location disclosure level. As participants change the disclosure level, IMBuddy provides dynamic feedback showing them how their changes would affect the information that would be disclosed to AIM buddies in that particular group (Figure 14, right).

The second interview session with the participants followed after 2 weeks of continued usage of the IMBuddy system. At this session, each participant reviewed their own location disclosure history. This access log provided participants with a history of every location request that was sent to IMBuddy (i.e., the screen name who sent the "whereis X" request to imbuddy411) and shows the level of location information that was received by the asker. After viewing the disclosure log, participants were given an opportunity to reflect whether they felt that: (1) the disclosed location information was inappropriate (i.e., too much location information was disclosed) and (2) their initial disclosure settings needed to be changed.

In the last week of the user study, a "stalker-bot" called "jasonkats722" was introduced. The stalker-bot was implemented as an AIMbot that would randomly request location information for each of our participants, two to three times per day. The stalker-bot was deployed near the end of the study, giving participants enough time to become familiar with how IMBuddy works and also enough time to settle into a "comfortable" disclosure configuration. Introducing the stalker-bot also ensures that the disclosure setting for the 'default' group is sufficiently and equally tested for all participants since "jasonkats722" is guaranteed not to appear in any of their manually defined privacy rules.

At the end of the study, a third and final interview session is conducted with each participant. At this session, participants completed a Likert-style questionnaire, asking them to rate their overall privacy comfort levels with their disclosure configuration. And, like the second session, participants were again asked to review their disclosure history. The intention here is to discern whether participants were aware of the stalker-bot and if they were comfortable with the location information that was shared with it. It is important to note that the stalker-bot was introduced into the study without informing the participants that the bot was in fact an artificial entity. At the end of the final interview session, this manipulation was revealed to all of the participants.

Figure 14. (left) Group-based privacy settings for disclosing location information using IMBuddy; (right) Disclosure History Page showing who asked for the user's location, what location information was shared, and at what time the location information was disclosed.

5.4   Results

IMBuddy was deployed to 15 students for four weeks. These participants were all medium to heavy IM users, as they had an average buddy list size of 120 screen names and averaged 1580 instant messages (including both incoming and outgoing messages) per week. These users also all scored as being "privacy pragmatic" according to the Westin privacy scale [Westin, 1991]. Across the four-week period, imbuddy411 made

175 location disclosures. There were 61 distinct screen names who queried IMBuddy and 15 of those were repeat requestors.

All participants considered location information to be potentially sensitive information and agreed that they do not carelessly disclose their location to others (μ=4.1, σ=1.1, on a 5-point Likert scale). When configuring their default disclosure settings, ten of the 15 participants used location abstractions (Figure 15). Recall that this disclosure level revealed whether users are on or off-campus, and included city-level information if they are off-campus. Three of the 15 participants chose to disclose no location information by default and two participants chose to disclose the most detailed location information as their default disclosure policy. Participants agreed that IMBuddy's location abstractions were easy to understand (μ=4.4, σ=0.5, on a 5-point Likert scale).

Figure 15. Distribution of participants according to what disclosure level they choose as their default disclosure policy for sharing their location information using IMBuddy.

After both the second and third sessions, none of the participants opted to change their initial disclosure settings. After viewing the disclosure log in the second session, participants could not recall any inappropriate disclosures that were made and were comfortable with the location information that had been disclosed to others.

At the first interview session, participants reported that they were comfortable with their default disclosure settings (μ=4.0, σ=0.9, using 5-point Likert scale). At the last interview

sessions, after the stalker-bot had been introduced for about a week, users' comfort levels were not significantly different from their initial self-reports.

12 out of the 15 participants noticed the stalker-bot prior to attending the last session. Most participants reasoned that jasonkats722 was perhaps one of their buddies, or that he was an old friend that was no longer on their buddy list. Most importantly, none of the participants were concerned that the stalker-bot had requested their location information, as they were all comfortable and confident in their default location disclosure settings.

5.5   Discussion

The main contribution of the IMBuddy study is that it provides empirical evidence that: 1) many users prefer and actually do choose location abstractions as their default disclosure policy, and 2) users are comfortable using location abstractions, even after a four-week deployment of a location-aware application and the introduction of a stalker-bot requesting their location information.

The majority of the participants (10/15; 66.7%) opted to share location abstractions as their default level of location information. Choosing this setting for the default group is significant because this group is most likely to contain individuals who the participant is not as familiar with. In other words, the default group is most likely to contain individuals with whom the participant has weak social ties. And, as previously explained in Section 2.5, providing social awareness for weak ties can help strengthen that social relationship. Thus, the observation that participants are comfortable with sharing location abstractions with this type of group provides promising evidence that disclosure abstractions can serve as a privacy-sensitive compromise for sharing location information in social applications.

However, there are at least two important limitations to the IMBuddy user study that are worth noting. First, it is not clear what would have occurred had there only been two

disclosure options given to participants. IMBuddy provides three disclosure levels, but if it were to use an all-or-nothing disclosure model (where participants must choose between disclosing no location information or disclosing fully precise location information), then it is unclear whether the 10 participants who chose location abstractions would end up joining the three people who chose to disclose nothing, or the two people who chose to disclose everything. Ideally, the introduction of location abstractions would encourage previously "disclosure-shy" participants to be more willing to share their location information. But, even if location abstractions may take away from the disclose-all category, it may still be an important feature if it provides evident that those users are more comfortable with disclosing location abstractions than with disclosing their precise location. However, this comparison was not done in the IMBuddy study.

Another limitation of this study is the relatively low usage of the IMBuddy system. A total of 175 location disclosures over a four-week period translates into an average of 0.73 location disclosures per day per participant. Thus, most participants probably only shared their location information at most once or twice each day. There were also very few location requestors. With 15 consistent screen names using IMBuddy to request location information for all 15 participants, it is likely that each participant had only one user, or two at most, asking for their location. Furthermore, with a laptop-based IMBuddy client, the range of potential locations that could be share was also limited. People are unlikely to carry their laptops with them everywhere they go and, even if they do, it is unlikely to be turned on so that IMBuddy's WiFi triangulation can compute the user's current location. Thus, it is possible that these three factors resulted in participants feeling little to no privacy threats in terms of revealing sensitive location information. Had participants been exposed to more location disclosures, it is possible that their self-reported privacy comfort levels might be lower.

5.6   Summary

In this chapter, we presented results from our second user study, which examined how location abstractions are utilized in end-user privacy configurations. While there has been significant prior work in this area, the main distinction of our work here is that it provided an *empirical* evaluation of a specific style of privacy controls (group-based) that used both geographic and semantic abstractions in a real-world, deployed location-aware social application. Past work has focused on various styles of privacy controls, though mostly in a controlled laboratory setting. Results from our one-month field deployment of IMBuddy suggested that users were comfortable with sharing *geographic* location abstractions as their *default* sharing preference *across all relationship types,* ranging from strong social ties (e.g., family and close friends) as well as relationships that had weak or no social ties (e.g., acquaintances or strangers) to the user.

The limitation of this result is that, based on this study alone, we are unable to determine whether providing geographic abstractions actually caused more people to re-consider sharing *none* of their location information, as opposed to sharing *some* of their location information. Our intuition is that the imprecision inherent in location abstractions did enable users to be more comfortable with sharing that level of location information, but this preference still needs to be empirically validated. To address this issue, we designed a second user study (described in the next chapter) to further explore the implications of supporting location abstractions in privacy controls, but with a specific focus on drawing out the differences between LSA privacy policies that do include location abstractions versus those that do not include such abstractions.

6. Privacy Configurations, Part 2

In this chapter, we continue exploring our second research question, which is to examine how location abstractions impact end-user privacy configurations (Figure 16). However, in contrast to the second user study, for our third study, we are examining a different type of LSA (one that synchronously shares one's location *history* vs. just one's latest location information). We are also interested in empirically determining if there are any differences between privacy policies when users are presented with LSAs that do offer abstractions versus those that do not offer any abstractions (i.e., the standard default style of social location sharing seen in many commercial LSAs).

reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 16. The four research questions covered in this dissertation. This chapter focuses on the second research question and third user study, which looks at how location abstractions are utilized in privacy configurations. In this chapter, we examine privacy policies for an LSA that synchronously shares one's location history (as opposed to one's current location information). We use this new style of location sharing application as a framework for conducting a comparative study between end-user privacy configurations that do and do not include location abstractions.

6.1   Motivation

In the previous chapter, we examined users' privacy configurations in order to better understand how they felt about sharing location abstractions using a realistic location sharing application. However, it is important to note that IMBuddy's embodies a specific style of social location sharing, namely that it only shares a user's current location.

Recently, several location sharing applications have begun integrating with online social network sites (SNSs), like Facebook, in order to leverage the much larger networks that these sites possess. Many SNSs are also beginning to add their own location sharing [Bilton, 2010; Siegler, 2010; Twitter, 2010]. With location sharing integrated into SNSs' content feeds, what was once just sharing of one's *current* location information, now becomes sharing of one's location *history*. Current examples of LSA-linked SNSs (Foursquare) and location-aware SNSs (Twitter, Facebook) rely solely on users manually reporting their location. However, it is inevitable that these services will soon support automated location sharing, especially since both the hardware (e.g., GPS-enabled phones) and software infrastructure for this type of reporting is already in place today.

Introducing automated location disclosures in SNS has significant privacy implications for users. Past work has examined privacy controls only within the context of sharing one's *current* location (e.g., [Consolvo, Smith et al., 2005]). As the gap between LSAs and SNSs continues to shrink, it is important that we also consider the scenario of sharing one's *past* locations. To this end, we defined another location sharing application, called Social Beacon, that focused on socially sharing one's location history with others. Using Social Beacon, we re-evaluated how location abstractions impact disclosure configurations. Specifically, we are interested in seeing whether the evidence found in IMBuddy (that supports using location abstractions) generalizes to a disclosure protocol that facilitates sharing location trails, instead of just single location instances. The two research questions that we focus on are:

- What kinds of privacy rules are created when users consider sharing their location history with others?

- How do simple manipulations, like offering different disclosure options, change one's privacy rules?

6.2   User Study

We conducted individual interviews with 30 university participants. Participants were recruited using a university-wide mailing list and pre-screened to exclude those with prior experience using LSAs. We opted to study novice LSA users to remove any possible interaction between privacy preferences and prior location sharing practices. Participants were 20-54 years old (μ=28.1, σ=7.3); 18 were male (60%). 10 participants were undergraduates, 11 were graduate students, and the remaining 9 were staff members. Of the 21 students, 13 (61.9%) had non-engineering majors. 20 participants reported using SNSs $\geq$ 3 times a week.

At the start of the interview, we introduced participants to "Social Beacon", a new location-aware SNS. In reality, Social Beacon was only a hypothetical LSA we used in order to provide a realistic grounding for probing privacy concerns. To ensure its realism, we pre-screened for active smartphone users, as such phones would have been the type of platform that Social Beacon would be deployed on.

Each hour-long interview began by introducing Social Beacon's sharing features, which allowed users to share their current and past locations with other Social Beacon users. Participants read several user scenarios for Social Beacon using relatively polished screenshots of the mobile interface for Social Beacon (Figure 17). These scenarios were advertised to participants as a way to increase their awareness about others (e.g., find out where friends went on their last trip or where they went last night), meet new friends (e.g., find others who frequent the same places), and get place recommendations (e.g., based on past visits, someone recommends a new restaurant). Participants were told that

Social Beacon shares, by default, a precise geographic description of their locations (Table 4) whenever another Social Beacon user selects their name.

Figure 17. Four scenarios shown to participants, during their introduction of Social Beacon's location-aware features: (left to right) finding friends who are near you, browsing situated status updates, viewing hotspots based on friends' past locations, finding people who were at the same place as you

	Types of Location Abstraction	Example
baseline condition	specific geographic description	street address or cross-streets/intersection
experimental condition	specific geographic description	street address or cross-streets/intersection
	general geographic description	city or neighborhood
	specific semantic description	business names like "Starbucks" or personal labels like "home", "work"
	general semantic description	types of places like "coffee shop", "restaurant"

Table 4. List of location abstractions used in Social Beacon and examples for each type. In this study, we compared two different privacy configurations. In the baseline condition, users were only given the choice of disclosing nothing or a specific geographic description of their location. In the experimental condition, users had the option of choosing from three types of location abstractions: a specific geographic description, a general geographic description, a specific semantic description, or a general semantic description.

6.2.1 Privacy Configuration Exercises

In the second part of the interview, participants defined privacy rules to specify who, what, and when they would share location history with others using Social Beacon. We conducted a within-subjects study with two configuration styles, both presented as paper-based exercises. We felt that the low fidelity nature of these exercises would give participants more opportunities to openly express their privacy preferences with a think-aloud protocol. To maintain the realism of Social Beacon, participants were told that, given the system's complexity, the experimenter would help users transfer whatever privacy settings they generated from the paper-based configuration exercises.

The two privacy configurations differ only on what type on location descriptions can be shared with others (Table 4). The baseline condition borrows from the "all-or-nothing" privacy settings that many LSAs use today for sharing current locations. The experimental approach gives users three choices of location abstractions that vary both the type of description (semantic vs. geographic) and the level of precision (general vs. specific). These four types of location abstractions were selected to broadly cover the range of descriptiveness in between the two extremes of the "all-or-nothing" disclosure model. These abstraction categories are also loosely based on the taxonomies from prior work as well [Lin, Xiang, Hong, and Sadeh, 2010].

In both configuration styles, participants could add subordinate conjunctions (e.g., "except", "only if") using four disclosure variables (Table 5). Past work has shown these variables are often used to frame privacy decisions for sharing *current* locations. We included them in our configurations to see whether they are also influential for sharing *past* locations too. To help ground these variables for Social Beacon, participants were told that variables like mood would be determined based on periodic self-reports. For example, Social Beacon would assume participants are in a good mood, unless they specifically tell Social Beacon otherwise (and the default setting would reset the next day).

Type of location sharing filter	Example (only share my location if …)
Time	… it's 5pm-8pm or …it's in the morning
Day	… it's during the weekend
Frequency	… I've visited this place at least 2 times
Movement	… I'm currently driving or walking
Mood	… I'm in a good mood

Table 5. List of filters (and examples of each type) that participants were allowed to add in order to express their sharing preferences. These filters were chosen based on past user studies that have examined how users decide what location information to share (though typically for purpose-driven location sharing scenarios).

Both configurations required privacy rules to be defined for specific relationship types. All the participants used the same relationships: strangers, classmates & coworkers, bosses & professors, acquaintances, casual friends, close friends, spouse or significant other, and family members. These types were borrowed from past work on location sharing [Anthony, Kotz, and Henderson, 2007; Consolvo, Smith et al., 2005]. To ensure the configurations were realistic, participants were told that they would later populate these groups by listing specific individuals in a separate interface.

The order for presenting the two configurations was counter-balanced. At the end of the interview, a post-study questionnaire was given and we revealed our experimental manipulation that Social Beacon was not a real system.

6.3   Key Findings

6.3.1   Rule Features for Sharing Location History

Our participants created 121 and 145 privacy rules in the baseline and experimental conditions, respectively.

Table 6 and Table 7 show how many participants shared their locations for each relationship type in both conditions. As expected, in the baseline condition, participants are more willing to share with spouses & family and least willing to share with bosses & strangers. When considering all relationships, participants were more likely to refrain from sharing in the baseline condition; in the experimental approach, participants were more likely to share general geographic descriptions.

		No Location	**Specific Geographic** *(address or intersection)*
A	Spouse/Sig. Other	3.3%	96.7%
B	Family	33.3%	66.7%
C	Close Friends	66.7%	33.3%
D	Casual Friends	96.7%	3.3%
E	Acquaintances	96.7%	3.3%
F	Classmates/Coworkers	93.3%	6.7%
G	Boss/Professors	100.0%	0.0%
H	Strangers	100.0%	0.0%

Table 6. Percent of participants that shared their past locations in the baseline condition, sorted by relationship type and type of location abstraction.

		No Location	**General Semantic** *(categories like coffee shop, restaurant)*	**Specific Semantic** *(business names like Starbucks or labels like "home")*	**General Geographic** *(city or neighborhood)*	**Specific Geographic** *(address or intersection)*
A	Spouse/Sig. Other	0.0%	63.3%	10.0 %	10.0%	53.3%
B	Family	16.7%	20.0%	0.0%	66.7%	16.7%
C	Close Friends	2.2%	13.3%	53.3%	30.0%	30.0%
D	Casual Friends	40.0%	16.7%	0.0%	53.3%	0.0%
E	Acquaintances	53.3%	6.7%	0.0%	43.3%	0.0%
F	Classmates/Coworkers	30.0%	13.3%	0.0%	70.0%	0.0%
G	Boss/Professors	46.7%	0.0%	0.0%	53.3%	0.0%
H	Strangers	63.3%	0.0%	0.0%	36.7%	0.0%

Table 7. Percent of participants that shared their past locations in the experimental condition, sorted by relationship type and type of location abstraction.

Table 8 shows how often each disclosure variable was referenced in participants' privacy rules. While there are several references to mood, there were also many references to time-based variables. This suggests that location-aware SNSs should considering

incorporating time variables into their privacy controls, particularly since time (as opposed to mood) is relatively easy to capture.

	Time	Day	Frequency	Movement	Mood	Time+Day	Time+Day+Mood	None
B	5	11	0	5	9	8	17	66
E	8	11	2	5	8	4	3	107

Table 8. Number of rules that referenced at least one filter (i.e., disclosure variable). A rule can contain multiple variables or no variables. The top row is for the baseline condition; the bottom row is for the experimental condition.

We also found that, in the baseline condition, 23.6% of the rules used negative sharing language (i.e., "do not share my location if I'm in a bad mood"). In contrast, this occurred in only 10.8% of the rules in the experimental condition. While these instances form only a portion of all the privacy rules, it does suggest that privacy configurations should consider including negative phrasing of rules, as it may be easier for participants to define their sharing preferences that way. Current LSAs rely on a white-list approach using positive sharing language (i.e., "only share my location under conditions X"), so supporting negative phrasing would likely require non-trivial architectural changes.

6.3.2   Comparing Configuration Approaches

Past work has shown SNSs typically have networks with more weak ties (e.g., casual friends) than strong ties (e.g., close friends) [Ellison, Steinfield, and Lampe, 2007]. This property has important implications for sharing location history. When, considering only weak tie relationships (i.e., rows D-F in Table 1, top), we see that in the baseline approach very few participants shared their location history (4.4%). In the experimental approach, more than half shared their locations (67.8%). This difference suggests that offering additional location granularities can lead to big differences in sharing. Though it may not be the most precise description, offering additional granularities can significantly encourage participants to share at least some of their past locations when they would not have done so in the baseline approach (t=8.38, p<0.001).

Another important difference we found is that the experimental approach resulted in much simpler privacy rules. First, fewer privacy rules had subjective conjunctions (z=3.02, p<0.01). Second, fewer privacy rules referenced >1 disclosure variables (z=4.02, p<0.01). This meant participants were more likely to share rules like "share my general geographic location always" (no subjective conjunction) vs. "only share my location if I'm in a good mood and it's a weekend" (multiple variables). We also found significantly fewer references to the mood variable in the experimental approach (z=3.33, p<0.01), which is noteworthy as accurately capturing mood is difficult to do. These results are also promising since having simpler rules for one type of sharing (location history sharing) will hopefully provide a more scalable solution when SNSs start including other types of context sharing and privacy rules.

Participants reported higher comfort levels on a 5-point Likert scale (μ_e=3.9, σ_e=0.68; μ_b=3.2, σ_b=0.71) when using the experimental approach (t=4.82, p<0.001). Given that participants shared more past locations in the experimental condition, it is encouraging that their more open sharing preferences did not adversely affect their comfort levels.

6.3.3   Reframing Privacy Configurations for Location Sharing

Past work has shown that the decision to share one's current location is often based on who sends the request [Consolvo, Smith et al., 2005; Lederer, Mankoff, and Dey, 2003]. In current LSAs that use our baseline approach for sharing current locations, this identity-centric strategy often results in infrequent location sharing and/or sharing with only a small number of friends [Moore, 2010]. To avoid this scenario when SNSs shift to sharing location history, we explored the effects of reframing privacy configurations to consider both *who* and *what* information should be shared. Our results show that providing relatively simple granularity options can significantly change how privacy rules are defined and under which conditions location information is shared. In particular, we provide empirical evidence that sharing more abstract location information (e.g.,

general semantic or general geographic descriptions) can lead to users sharing their location history with more people in their social networks. This provides a nice compromise for service providers in that they have more users engaging in location sharing while also providing users with some ambiguity and plausible deniability with the abstractions.

It is worth noting how our results compare with Foursquare, a LSA that shares specific semantic and geographic descriptions of users' current locations. Foursquare lets users link their accounts to Facebook to encourage more location sharing; however only 28% currently do this [Moore, 2010]. Our results may provide one explanation why this is a somewhat under-utilized feature. The disclosure pattern that Foursquare supports was chosen by relatively few of our participants and was only chosen when sharing with close friends (16.7%). When sharing with weak ties, participants preferred general semantic and geographic descriptions, a combination that Foursquare does not support yet. This result provides further evidence that LSAs should consider supporting broader definitions of location disclosures, particularly if they want to encourage automated (vs. selective) disclosure of location history in SNSs.

In conclusion, we interviewed 30 participants and asked them to completed two privacy configuration exercises for specifying their preferences for sharing their past locations with others. The goal for our comparative study was to provide a more controlled evaluation of how disclosure abstractions can influence end-user privacy concerns at the privacy configuration stage. Our results suggest that offering certain simple location abstractions (i.e., general geographic descriptions and general semantic geographic descriptions) can result in privacy rules that have fewer exceptions. In addition, these abstractions are likely to make users feel more comfortable sharing more of their location history with others.

6.4   Summary

In this chapter, we presented results from our third user study that was designed to continue our exploration of how location abstractions are utilized in end-user privacy configurations. One important difference is that, unlike the second user study, this study examined sharing of *past location* information, as opposed to only sharing current location information. This meant that users were sharing multiple locations for each incoming request for their location information. Designing a LSA like this introduces significantly more privacy concerns for the end-user as they are now sharing much more information than they previously were (i.e., when using an LSA like that introduced in Chapter 5. Thus, we wanted to see what kinds of sharing preferences users had when exposed to this relatively new type of LSA. We found that our participants preferred to create privacy rules that reference *general semantic* and *general geographic* location abstractions. In the previous chapter, the preference was for *geographic* location abstractions. Thus, one potential reason that there is a different abstraction preference between the two studies could be attributed to the fact that the LSA used in each study supported different types of location sharing.

Another important aspect of this study that was covered in this chapter was the comparison between a privacy configuration style that included location abstractions and one that did not (and only offered the option to either not share anything or to share a fully precise location description). The motivation for this conducting such a comparison was based on results from our second user study (Chapter 5). Thus, our focus was to look at the additive value that location abstraction provides for end-user privacy configurations. Our results indicated that, by offering additional disclosure options in the form of location abstractions, users were more likely to share their location information with a wider audience. We also found that the resulting *rule-based privacy configurations* were *simpler* (fewer overall rules) and *less complex* (fewer caveats and exception clauses)

when users were given the option to include location abstractions in their rules (in comparison to a configuration interface that did not allow any use of abstractions).

7. Sometimes Less is More: SLIM Visuals

In this chapter, we move on to our third research question, which is to examine how visual representations of location abstractions can influence end-user privacy concerns and sharing preferences (Figure 18). So far we have discussed how abstractions impact people's decision making and their privacy configurations for social location sharing. However, these two events typically take place prior to the actual exchange of location information.

In our third research question, we are interested in examining users' privacy concerns during location disclosures. Typically, a large part of what occurs when information is exchanged between users is that they are presented with visualizations of other people's location information. Our insight into this scenario is that it is possible for different types

reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 18. The four research questions covered in this dissertation. This chapter focuses on the third research question, which examines how differences in location visualizations can impact end-user privacy preferences and their perceived utility metrics for social location sharing.

of location visualization to adversely impact users' privacy preferences. In particular, by adding location abstractions to LSA, we are enabling more types of information to be shared visually. Thus, we wanted to explore different ways that abstractions could be visualized and to see whether adding an additional layer of complexity is worth implementing in future LSAs.

7.1   Motivation

In Chapters 5 and 6, the primary focus was on examining how location abstractions influenced users' privacy configurations for location sharing. We examined sharing preferences in two types of social location sharing applications: one that shares a user's current location (IMBuddy), and another that shares a user's location history (Social Beacon). The studies for both of these location applications were designed to introduce realistic privacy threats for users. In IMBuddy, users have no control when others may ask for their location information. In this LSA, location information is disclosed whenever someone asks for a user's information. Though the descriptiveness of the information can vary (according to the user's privacy preferences), disclosure are request-driven in IMBuddy. In Social Beacon, users also have no control when their information is shared, because again location sharing is initiated by the requester, not be the discloser. In additional, each location request in Social Beacon results in others being able to view up to one week of the user's past locations (and not just one's current location, as in IMBuddy). Thus, in both systems, there is a clear incentive for users to provide "good" default privacy settings that adequately match their actual sharing preferences.

Thus far, the assumption has been that a user's perceived privacy concerns are largely driven by their own mental models for how their location information should be exchanged with others. While this certainly plays a large role (and is evident in our qualitative interview findings from our second and third studies), there are also other ways that users can be influenced to have different privacy preferences (which can in turn affect their location sharing behaviors). We hypothesize that one of those ways is through

location visualizations. In particular, if LSAs depict information differently, then we anticipate that end-users will have different reactions to their information being shared. Thus, an application's visualization component is arguably an important aspect to whether an LSA will be warmly embraced by its user. Better visualizations can directly impact the (perceived) usefulness of information sharing, as well as influence users' perception of how much data is being shared about them. In this study, we intend to begin the exploration into this dimension of LSA by conducting a study that broadly looks at end-user perceptions of different types of visualizations for LSAs that record and share a user's location history with others.

7.2   Pilot Study

As part of the Social Beacon study (described in Chapter 6), we also probed participants as an initial probe of whether visual representations of (the same) location information can influence users' perceptions in terms how comfortable they are with sharing their location information with others. In the pilot study, we selected two common visualizations in use by current commercial location sharing applications: 1) a map-based visualization (akin to what is used by LSAs like Google Latitude [2009]), and 2) a text-based visualization (similar to Twitter's geolocation feature [2009b] and what Facebook's Places [2010] supports).

We created hypothetical screenshots for these two visualizations (Figure 19). Each visualization contained location trails from three individuals (i.e., the user's friends). The scenario presented the user is that, by initializing the history-sharing LSA, one of two visualizations would appear. The map-based visualization showed the three sets of location trails, differentiated by their different colored markers. Each set of markers represented a particular friend's location trail. We varied two marker variables: its size and its transparency. We told participants that the size of the markers represented roughly how much time was spent in that location. Larger markers meant that their friend spent more time at that place. Participants were told that more transparent markers represented that their friend visited that place further in the place. In other words, the most

transparent marker represented the first place that the LSA recorded for their friend's location history. The least transparent marker meant that it was the most recent place that the LSA had recorded for that friend.

The text-based visualization showed a sequential listing of each friend's information, separated by three different tabs. Each tab showed the same historical information: a location label, the time that the friend arrived at that location, and the length of time that the friend spent at that location. Tabs were arranged with the person whose location was updated most recently.

These two visualizations were designed to be "content-equivalent". Specifically, the location information that is present in the map-based visualization is also present in the text-based visualization. Furthermore, one could presumably convert the text-based visualization into the map visualization (and vice versa) without any additional information. For example, clicking on the markers in the map-based visualization revealed the times that the friend visited that particular place and precisely how long they had stayed at that place; thus, both visualizations contained arrival and duration information. It is also possible to reverse geocode each marker's geographical coordinates. In addition, the user can click on the markers in the map-based visualization to see additional location labels (like the labels used in the text-based visualization). Thus, both visualizations have the same spatial location information as well. So, information-wise, the map-based visualization and text-based visualization are content-equivalent and one can easily translate the information from one visualization into the other.
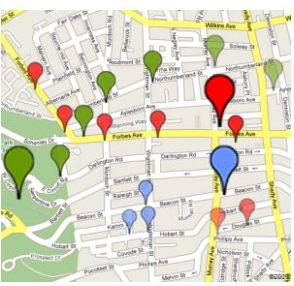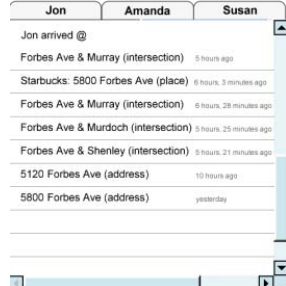
	map-based	text-based
"I prefer to see my friends with this view"	23	7
"I prefer to show up on this view"	11	19

Figure 19. Number of responses from participants when asked which visualization they would: a) prefer to view their friends' past locations and b) prefer to display their own past locations to others. The map-based visualization shows a spatial representation with three sets of colored markers (green, blue, and red), each representing a particular friends' location trail. The text-based visualization shows a sequential listing of each friends' information, separated by three different tabs.

We asked thirty participants which visualization they would: a) prefer to use when viewing their friends' past locations and b) prefer to appear on when their friends check to see their own past locations. Interestingly, we found that there was a clear preference for users wanting to use the map-based visualization to view their friends (76.7%). When considering which visualization they would like to appear on, their preferences shift to the text-based visualization (63.3%).

When asked to explain their visualization preferences, many participants mentioned that the map-based visualization seems like it provided more information, or the information provided was much easier and much quicker to parse when presented within a map:

> P3: *"Looking at the map – I can see more information. The extra information might come in handy…you never know. It's nice to have just in case…"*

> P11: *"It's easier to see everything. I can just glance at it and roughly know what's going on...the other one just seems too...tedious. It's just kind of hard for me to read through quickly."*
>
> P27: *"I don't like how I have to click through to see all my friends in the [text] one. The [map] one is quicker – I don't have to mess with the screen. Everything is just there."*

Other participants reported that the map-based visualization simply seemed more familiar to them and more conducive for physically finding the person (whether for functional purposes, like coordinating a meeting, or just for social purposes):

> P16: *"I guess it's just because I associate location stuff with maps. It just makes more sense to me that if I want to where someone is, then I should look on a map, you know?"*
>
> P19: *"It's more useful to see people on a map. If I decide I want to meet them some place, I can sort of figure out how to do that, you know, by just looking. On the [text] one, I have to think about...at least a little bit. It's just more difficult....at least for me..."*

When explaining their choice for which visualization they would prefer to appear on, participants cited concerns about their fears regarding location tracking.

> P6: *"Personally, the map seems like it could get out of hand. Someone could keep tracking my location, right? And they would know where I've been? I dunno – sometimes it's kind of nice if people don't know where you are."*
>
> P14: *"The text one just seems like really dense to me. Like I don't even want to read it. So I guess if other people are like me, then if I pick it [the text visualization], then they probably won't bother looking me up."*

The different preferences participants claimed for how they wanted their location visualized reinforces that there are important privacy considerations for LSAs that go beyond designing adequate privacy controls and feedback mechanisms. Based on

feedback from the participants, our results suggest that the presentation of location information can indeed influence users' perceived privacy comfort level.

As a follow-up to this pilot study, we conducted a second, more rigorous study that focused on a more methodical evaluation of users' visualization preferences for location sharing. In the next few sections, we describe our specific research questions, our study protocol, and our results.

7.3   Research Questions

In order to provide a more systematic exploration, we designed a comparative study to better understand the interplay between visual representations of location information and end-user sharing preferences. Specifically, we are looking to address the following research questions:

- How do different visual presentations of location information affect users' perceived privacy, in terms of their willingness to share their location history?
- What is the best way (from an end-user perspective) to incorporate location abstractions into these visualizations?
- What specific visual elements make one type of location visualization more (or less) "acceptable" to a user, as measured by their willingness to share their location information with others?

The last research question is particularly important. When considering the location applications that support sharing of location history, there are many variables that could be visually emphasized. If a particular design emphasizes the wrong location variables, then it is entirely possible that the LSA will alienate users, resulting in them feeling uncomfortable about sharing their location history information with others. With the results of our study, we hope to provide initial insights about which location variables should be avoided and whether certain visualization styles can lead to users feeling more comfortable about the privacy and utility aspects of location sharing.

7.4   Examining the Dimensions of Sharing Location History

The first-generation of social location sharing applications were designed primarily to share a user's *current* location. Sharing this data goes beyond just revealing the geographical coordinates for that location (i.e., a location's spatial information); it can also include sharing information like: when the user arrived at a location, how long they've been at a location, and what is an appropriate label for name for a location (e.g., is it a place of business, like "Starbucks", or a personal place, like "home"?).

However, when we considering location applications that share *past* locations, there are many more variables at the application's disposal which could be shared with others, including:

- the sequence of locations (i.e., in what order did the user visit each place?)
- the arrival time for each location
- the departure time for each location
- the frequency of visits (i.e., how often does a user visit a particular place?)
- the total time spent at a place (accumulated over all visits or for each visit)

These variables are, of course, in addition to the variables that can be shared when considering LSAs that only support sharing of current location information.

In the pilot study, the map-based visualization that we used only emphasized a subset of these variables. Specifically, the visualization shared three aspects of a user's location history. First, placing the marker on the map inherently conveys the spatial position (i.e., the geographical coordinates). Second, markers of the same color were shown at different levels of transparencies, which indicated how recent the user had last visited that place. The most opaque marker corresponded to the most recent location that the user had visited. Third, markers of the same color were shown at different sizes; this corresponded with how much time the user had spent at a particular location (accumulated over all visits). Locations with large markers thus indicated that a person had stayed at that location for long periods of time. We limited our selection to only three variables, as we

felt that this amount of information would not be too overwhelming for users and would still provide enough details about one's location history.

In our subsequent exploration of location visualization, we continue to emphasize these three aspects of location history and added the variable that looked at frequency (i.e., how often a person visited a place). We acknowledge that sharing location information also affords sharing of other variables, include the motion of the user (i.e., which direction is the user moving) and speed (i.e., how fast is the user moving). However, we felt that sharing these variables place less emphasis on physical locality and instead emphasize more the transitions that occur *between* places. While the difference is subtle, it is an important one. Sharing information about when people transition is more likely to emphasize more purpose-driven sharing, such as determining whether it is an appropriate time to call (e.g., is the person driving now?) or checking when someone will arrive at a meeting place (e.g., is the person on their way to the meeting place?). It is difficult to imagine scenarios where sharing transition information is useful for *social* purposes. In order to use location information for grounding purposes, it is more useful to share *places* that people visit, rather than the *times* that people are on their way to a place. Given this distinction, we omit visualizations that include these types of location variables. Not because we think they are not important; rather, given our dissertation's focus on *social* location sharing, we thought it would better to focus on the aforementioned location variables (i.e., a place's spatial coordinates, the arrival information, the duration information, and the sequential ordering).

7.5   Choosing Representative Visualization Styles

In the pilot study there was a two-way comparison between a map-based visualization and a text-based visualization. Preliminary results from the study suggested that one feature that participants found most appealing about the map-based visualization was that it was easy to glance at the graphics and immediately extract meaningful information from it. Based on this input, we decided to include a third visualization style that also

supported glanceability, but emphasized different location variables. This led us to choosing three distinct styles of location visualizations: text-based, map-based, and time-based.

7.5.1   Text-Based Visualizations

The text-based visualization we used in this study is based on the version used in the pilot study, with one major modification: each location report includes an explicit timestamp (see Figure 20). Though this timestamp can easily be computed based on knowing the current time and a relative time span (e.g., "5 minutes ago"), having the explicit timestamp listed helps users to quickly scroll through past places based using this information as a simple lookup index. Adding this feature also makes it clearer to participants that our location visualizations are indeed content-equivalent with each other.

The most salient features of the text-based visualization are: 1) locations are sequentially ordered, and 2) every location is treated the same (at least visually speaking). In other words, every single line in the visualization is visually no different than any other line. Unlike the other visualizations, the text-based visualization emphasizes a specific temporal dimension of location history: the *order* of the places that a user has visited. Thus, we can use this visualization to probe whether the sequential ordering of locations can potentially trigger end-user privacy concerns about location sharing.

Figure 20. Text-based visualization showing three location variables: when the user arrived, where a location is spatially oriented (in this case, it is described using a street address), and how long the user stayed there. Variables like how many times the user visited and the total time spent at a place need to be computed by hand when using the text-based visualization.

7.5.2   Map-Based Visualizations

Map-based visualizations are inherently saturated with information, especially in urban areas where map features like roads and highways can obscure other information being overlaid on top of the map. Thus, we wanted to add a minimal number of information layers to this type of visualization. In many location sharing applications, the most common way of indicating a set of locations is to use a marker (like a pushpin) than can "point" to specific geographical coordinates (Figure 21, left). This is in contrast to using a less precise marker (like a halo) that covers a larger geographical area (Figure 21, right).

For our map-based visualization, we chose to use the halo styled marker. The reason for this design decision is that it is more visually representative of the different types of location abstractions that we will introduce in a later subsection (7.5.4). The halo marker also provides more plausible deniability for users in that it covers a region of potential

places and users can easily fudge their true location (within the boundaries of the halo). In order to support this feature, we made sure that the user's true location was not marked by the halo's midpoint, as that would defeat the purpose of having a less precise marker. Instead, we randomly added noise δ (where δ ranged from +/- 300m) to the true location (latitude$_{true}$, longitude$_{true}$) so that the midpoint of the halo marker is defined by the GPS coordinate (latitude$_{true}$ + δ, longitude$_{true}$ + δ). The size of δ was chosen to reflect a blurring of up to three city blocks, which are on average about 100m long (per block).

Figure 21. Two different marker styles for map-based visualizations. The marker on the left "points" to specific geographical coordinates, whereas the marker on the right (a "halo") covers a much geographical area. The intention of including different marker styles is to explore whether the visual precision associated with location sharing can impact end-user privacy concerns.

In order to include the other relevant location variables (like those explicitly listed in the text-based visualization), we created interactive map-based visualizations. Whenever a user clicks on a marker, the corresponding location variables are shown, including: when the user arrived, when the user left, how long the user stayed, and a description of where the location is (in Figure 22, this is done by using a business name to describe the place).
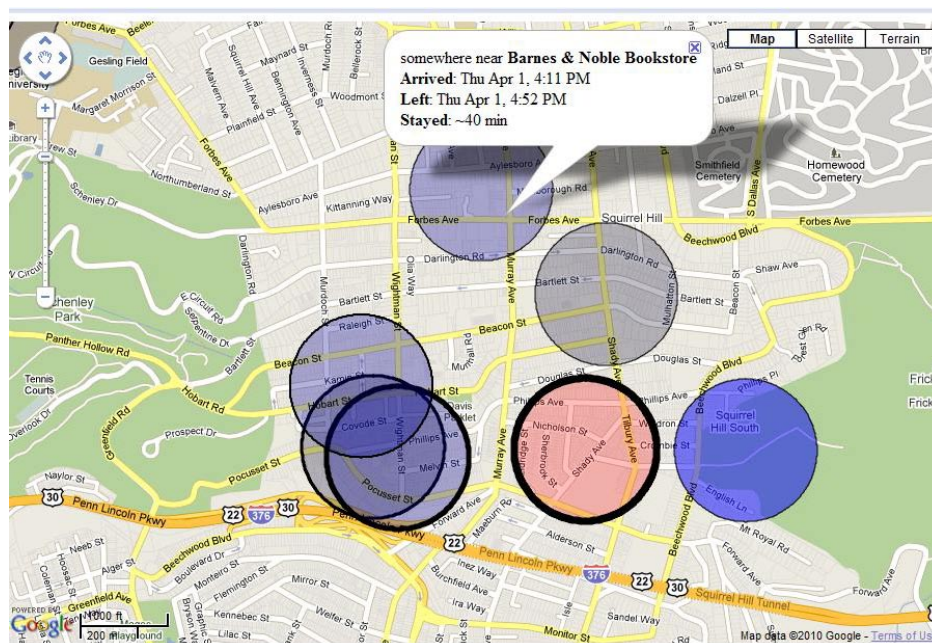
Figure 22. Map-based visualization that uses the halo marker and includes information about where the location is, when the user arrived & departed, and how long they stayed at the location. This particular visualization also uses semantic abstractions for naming the location (in this case, it uses a business name, aka a general semantic label).

Aside from the location label, we tried to also visually depict various temporal information relating to the user's location history, including when they arrived (i.e., the sequential ordering of the locations), how long they stayed at a place, and how often they visited a place. In theory, there are many ways one could represent these location variables. Card et al [1999] and Ware [1999] have proposed that, for every object, there are at least nine visual properties that can be manipulated to convey information, including the object's position, size, orientation, grayscale, color, texture, shape, animation, and transparency level. MacEachren proposed the use of other visual elements, such as resolution, crispness, and arrangement [1995]. Healey et al suggested that varying lighting, quantity, and depth could be useful for certain types of information visualizations as well [1995].

Based on the qualitative feedback we received in the pilot study, we found that users could easily interpret the meaning of larger vs. smaller marker sizes and understood that size correlated to time spent at a particular location. However, the difference between two transparency levels was not always easy for participants to detect. Most importantly, participants reported it was hard to consistently pick out the marker that was most opaque (which should correspond to the user's current location). Because of this we referred back to the basic visual elements mentioned by Card [1999] and Ware [1999] and opted to use a combination of color *and* transparency to indicate the sequence of location visits. In particular, the red colored halos indicate the most recent location that the user visited (i.e., where the user is currently located at, or where the user was last seen). All other markers are a different color (blue) and vary in transparency levels depending on how "stale" the location is. Note that this change also means that the map-based visualizations used in this study are different than the ones used in the pilot study. Here, the visualization only includes the location history for *one* user; the visualization used in the pilot study showed location trails for three friends' on the same map.

This left only one location variable that needed to be matched to a visual element, namely how often a user visits a location. We opted to convey this information by varying the width of a marker's border (see Figure 21, right). Thus, large markers with thick borders meant that the user visited a particular often and has accumulated a significant amount of time there. Small markers with thick borders meant that the user frequently visited a place, but never stays there very long. Varying the marker's border can be considered a crude approximation of varying a marker's texture, which is another one of Card and Ware's nine visual properties.

In programmatically developing these visualizations, we should note that the size of a marker (corresponding to how much time a user spends at a location) and the thickness of a marker's border (corresponding how often a user visits a location) is always computed relative to that particular user. In other words, what appears as a large marker for one user may appear as a smaller marker for another user, even though they may both represent the same amount of total time spent at a place. We felt that this design decision was

reasonable, particularly since each visualization now focused on only one user's location trail (as opposed to having multiple location trails for different friends). This disclosure model is similar to Facebook where, upon selecting a particular user, one is shown only that user's "wall", i.e., a newsfeed that shows a history of that user's past disclosure. Also, when considering social LSAs, these applications tend to be integrated with large online social networks. The average Facebook user has around 130 friends [Facebook, 2004b]; plotting this many users on a single visualization would be unwieldy and very difficult for users to quickly glance at.

In conclusion, the most salient feature for the map-based visualization is the spatial information of the user's location history. Though, there is a visual mapping between sequence and transparency, the order of the location visits is arguably less noticeable when compared to the text-based visualization. Instead, it is the markers and their placement on the map that is most noticeable. Thus, we can use the map-based visualization to probe how sensitive users are to sharing the spatial properties of their location history.

7.5.3   Time-Based Visualizations

In this study, we introduced a third visualization that also lends itself to glanceability: the time-based visualization (Figure 23). In this visualization, the emphasis is on the ordering of the locations. This is visually depicted using a timeline and color-coded blocks that correspond to when the user arrives and leaves a particular location. The colors of the blocks are randomly assigned, so the exact colors are not meaningful. However, blocks that are similarly colored (like the purple colored blocks in Figure 23) indicate that the user revisited a previous location. In these cases, the color of the block is selected to match the color of the first block corresponding to the same location.

In order to ensure that the time-based and map-based visualizations are content-equivalent (i.e., conveying the same location variables), we allowed users to interact with

the timeline in order to convey the non-temporal properties of their location history. Whenever the user clicks on a particular block, they see more details, including precisely when the user arrived & departed, how long the user stayed, a scaled down map showing the spatial orientation of the location, and a description for the place (the visualization in Figure 23 uses a street address to describe the location). Even though the time information is visually conveyed in the timeline, it was important to also provide the precise time information as well since the other two visualizations also provide this level of descriptiveness.
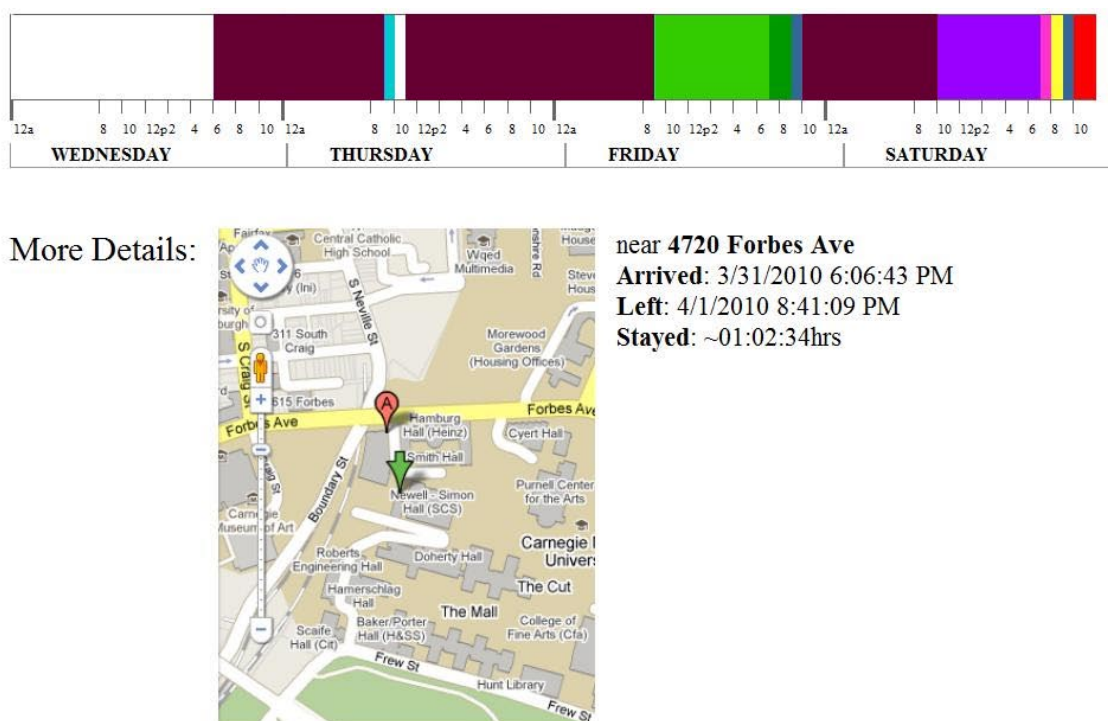
Figure 23. Time-based visualization showing four location variables: when the user arrived & left, how long the user stayed, where a location is spatially oriented and how to describe the location (in this case, using a street address).

Unlike the other visualizations, the most salient feature of the time-based visualizations are the color blocks shown on the timeline. These blocks emphasize how much time the user spends at each of their locations (because of the large size of the blocks). In addition, this visualization easily draws your attention to locations with repeated visits (because of

the similarly colored blocks). On the other hand, the spatial information is less emphasized in the time-based visualization, as user are required to first click on the colored blocks in order to see any map-related information.

7.5.4   Using Disclosure Abstractions as Location Labels

Thus far, we have not explicitly mentioned how location abstractions can be integrated in location visualization. Upon closer inspection of the three examples provided for the text-, map-, and time-based visualizations (Figure 20, Figure 22, Figure 23), one will notice that they each use different place labels to describe a location. These labels are based on different location abstraction and provide an additional dimension that we are interested in exploring in terms of visually representing location history to others. In particular, we want to explore how different location descriptions affect users' willingness to share certain location visualizations.

In this study, we explored four different location abstraction types. These are identical to the abstractions we studied in our first study (Chapter 4), which are:

- general geographic descriptions, such as city and neighborhood information
- specific geographic descriptions, such as a street address or intersection
- general semantic descriptions, such as the type of place ("coffee shop")
- specific semantic descriptions, such as a business name ("Starbucks")

Aside being a literal reference to the location being visually described (which applies to all three visualization types), these place labels place an additional role in map-based visualizations. In particular, the diameter of the halo markers are directly related to the type of location abstractions associated with the visualizations. In other words, different location abstractions will result in different sized halo markers. For halo markers associated with general geographic labels (e.g., city or neighborhood), we set the diameter to be two miles wide. Most of our participants' locations were for the Pittsburgh area and, given the city's layout, we felt that this was a reasonable distance that was representative of the conceptual size of local neighborhoods (e.g., areas like Shadyside,

Oakland, etc.) in the Pittsburgh area. For halo markers associated with semantic labels (e.g., "coffee shop", "Starbucks"), we set a smaller diameter (0.25 miles wide). Conceptually speaking, semantic labels are more precise than the general geographic description, so this is reflected in the smaller marker size. The exact size was somewhat arbitrarily determined, though the intention was to select a size that typically spanned several city blocks. For the major Pittsburgh neighborhoods, where most of our participants' locations were reported from, the 0.25 mile diameter seemed like a reasonable approximation for this.

7.5.5   Visualization Combinations

In this study, we are manipulating three different aspects of location visualization: the visualization type (text-, map-, or time-based), the marker type (a pointer type or a halo), and the location label (either a general or specific version of a geographic or semantic location abstraction). Combining these variables leads to twenty possible combinations (Table 9).

We presented eighteen of these visualizations for users to evaluate. We excluded two of the visualizations since they create an illogical combination. These two are the map- and time-based visualizations that use a halo marker with a specific geographic label (Table 9, the grayed out cells). The intention behind using the halo marker is that, by providing a large coverage area, users are afforded more plausible deniability in terms of being physically found. However, when you match the halo marker with a fully precise location description (e.g., a specific geographic label such as an address or intersection), the ambiguity provided by the halo becomes useless. Thus, we do not consider these two visualizations in our study and only present the other eighteen visualizations to our participants.

	Text-Based	Map-Based		Time-Based	
	(no markers)	pointer	halo	pointer	halo
specific geographic label					
general geographic label					
specific semantic label					
general semantic label					

Table 9. Shows the twenty different combinations of visualization type (text-, map-, or time-based) + marker type (pointer or halo) + abstraction type (specific or general, geographic or semantic). The highlighted cells are the eighteen visualization conditions that we evaluated in a within-subjects study. The grayed out cells are two visualization conditions that were not included in this study because they are not logical combinations. For example, it is meaningless to have a halo marker with a specific geographic place label, as any imprecision afforded by the halo marker is lost when using the precise geographic reference.

7.6 User Study

To ensure that participants realistically considered their privacy concerns when evaluating these location visualizations, we collected two weeks of actual GPS traces from twelve participants, all of whom were recruited through a university-wide mailing list. Participants ranged from 23-51 years old (μ=30.8, σ=6.2); five were female. Seven of the participants were graduate students; the remaining participants were university staff members. Participants were evenly split between those with prior training in technical (e.g., natural sciences, engineering) and those from non-technical fields.

7.6.1 Part 1: Entrance Survey

Participants completed a 10-min online survey to collect basic demographic and social network information. We intentionally did not ask include privacy to avoid biasing participants later in the study. For their social networks, participants provided examples

(names) for four relationships: family members, acquaintances, managers/bosses, and close friends. We told participants that their examples must live in the same city. This way we could control for geographical distance and avoid having that factor influence participants' visualization preferences.

7.6.2   Part 2: Location Data Collection

Participants were given mobile phones (Nokia N95s) to carry for two weeks and were required to use the N95s as their primary mobile phone during the study period. This helped to incentivize them to keep the phone sufficiently charged at all times, which in turn meant that we could continuously collect their location data.

The phones were equipped with location-logging software to collect participants' actual location traces (the same software used in [Benisch, Kelley et al., 2008]). The software ran continuously in the background (without user input), using both GPS and Wi-Fi positioning technology. To reduce power consumption, the application used the phone's accelerometer to selectively sample location information. The software we used was identical to what was used in our first user study (see Chapter 4.6.2 for more implementation details).

7.6.3   Part 3: Programmatically Generate Location Labels

Before each interview, we analyzed each participant's location trace. We used Skyhook's API [Skyhook Wireless, 2003] to translate Wi-Fi readings into GPS coordinates. We then computed the distance and speed between adjacent coordinates to determine if the participant was moving. Places that the participant stayed for more than five minutes were marked as "significant". Using the coordinates for each significant place, we programmatically determined the four different types of location labels.

To compute the general geographic description (i.e., city and neighborhood information), we queried a publicly available database to first obtain reverse geocode the geographical coordinates to a zip code and then used the zip code information to lookup the nearest neighborhood. To compute the specific geographic description (i.e., street address or nearest intersection), we used Geonames [2010] to perform reverse geocoding using their `findNearestAddress()` and `findNearestIntersection()` webservice methods.

To compute the semantic descriptions (i.e., the type of place such as "coffee shop" or "restaurant"), we used the specific geographic description (i.e., the street address or nearest intersection) to query the Google Maps API, which generates a list of the nearest places. Each of these results includes information about the type of place it is (e.g., "Restaurant", "Shopping", etc.) and the name of the place (typically a business name, like "Starbucks"). We record the top result, as it is supposedly closest to the place that we are trying to label. In order to generate additional label candidates, we also conducted similar lookup queries on several other publicly available database sources, including Microsoft's Mappoint webservice [2000] and Wikipedia [2001a]. To use Wikipedia, we first scraped Wikipedia for their geotagged articles and used these tags to create a local database of coordinates matched to location labels.

The challenge behind automatically generating location labels is that there is a good chance that the generated label is incorrect. In fact, there are several ways in which the label generation process is susceptible to errors.

1. Sensing Errors: All labels ultimately depend on obtaining accurate geographical coordinates. When the phone is able to lock on to a GPS signal, then these coordinates have relatively high accuracy. However, there are often times when GPS readings are not possible, e.g., when the user is indoors. In these cases, the phone relies on Wi-Fi readings. Switching between GPS and Wi-Fi sensing consumes a non-trivial amount of battery power. If the phone is not on, then it goes without saying that obtaining any type of GPS coordinates is impossible. While our software is designed to make it difficult for users to forcibly quit the

application, persistent participants (particularly those frustrated by poor battery performance) can find ways to kill the mobile application.

2. Triangulation Errors: For Wi-Fi readings, we rely on Skyhook's API to triangulate the data into GPS coordinates. This process is, by definition, only an approximation of the actual coordinates (i.e., ground truth). The accuracy of these coordinates is also highly dependent on how up-to-date Skyhook's database is. For a more detailed discussion of the shortcomings of Wi-Fi localization, we recommend the reader refers to any one of several Placelab papers (e.g., [Schilit, LaMarca et al., 2003]).

3. Interpolation Inaccuracies: Even with perfectly accurate GPS coordinates, there is still a reliance on public databases to provide accurate reverse geocoding services. However, by definition, reverse geocoding does not return actual addresses, only an approximation (i.e., a best estimate). For example, in order to determine the exact street number for a particular GPS coordinate, the reverse geocoding request often relies on some type of interpolation between two known (i.e., ground truth) street addresses. This interpolation process is, of course, not an exact science and even slight variation in GPS coordinates can result in drastically different reverse geocoding results.

4. Sparse and Stale Database: Assuming that we were able to retrieve a perfect GPS reading from the phone that was then perfectly interpolated into an address via reverse geocoding. We must then use this address to query a separate set of public database to find the nearest points of interests (e.g., restaurants, shopping malls, etc.), which we can then cull for location labels. Thus, we are entirely at the mercy of these services. Two problems that we have frequently encountered are: 1) these database contains out-of-date entries, and 2) the database only includes location information for relatively small geographic area. For example, for databases based on sources like Wikipedia, the amount of geolocation data is entirely driven by the generosity of users providing these tags. However, if a user ends up in an area that is not densely populated, then it is likely that these databases will not be seeded with enough information and, thus, no location labels would be generated.

7.6.4 Part 4: User Validation of Location Labels

We went through several series of data collection in an attempt to mitigate these errors. We ran at least four different trials of collecting 1-week long GPS traces for various pilot users. From these traces, we found that we could only consistently and reliably translate the recorded GPS sensor readings into geographic place labels. Semantic place labels proved to be much more difficult. As we were less interested in developing algorithms for intelligently guessing the most appropriate labels (e.g., through heuristics or machine learning techniques), we decided to modify our study protocol to include a human-in-the-loop mechanisms to verify and correct, if necessary, our programmatically generated place labels.

To do this, after a week's worth of data collection, we would post-process the GPS readings to extract the GPS coordinates of each participant's significant places. We would then programmatically run these coordinates through each of our database sources to generate the location abstraction descriptions, including the general geographic label, the general semantic label, and the specific semantic label. The specific semantic labels were then emailed to the participant, along with a timestamp of their stay at each of these places. We asked participants to verify and correct, if necessary, any obviously incorrect labels. Then, for any timestamps that were missing a location label (i.e., our automatic label generator was unable to find an appropriate match), we asked participants to provide several labels (on their own) to describe that place. To encourage both geographic and semantic labels, we provided several tutorial-like examples for participants to refer to. We then asked participants to consider if there were any other locations that should be added to the list. In these cases, the phone may have been off, resulting in there being no sensor readings and no way for the automatic label generator to even suggest location names. Because the study ran for two weeks, we repeated this process twice: once at the end of the first week of data collection, and again at the end of the study.

While this method certainly requires some effort on the part of the user, we felt that this was critical for the success of our study in order to fairly evaluate our location visualizations. In particular, we have found (through pilot studies) that when visualizations contain inaccurate labels, users have a tendency to judge visualizations based on their naming mistakes, rather than on any privacy-related preference. In order to eliminate this potential bias, we wanted to ensure that our visualizations were as accurate as possible, so that they would primarily then be judged by their visual properties (i.e., based on their properties of being text-, map-, or time-based).

7.6.5   Part 4: Evaluate Location Visualizations

We randomized the order in which we presented the visualizations to users, both in terms of the three main types of visualizations (text-, map-, or time-based) and in terms of the marker & label types. Because the label types are often not as visually noticeable at first glance, we made sure to highlight this particular difference between each of the visualizations. Otherwise, no other salient features were explicitly drawn out for our participants.

For each relationship type (family members, acquaintances, managers/bosses, and close friends), we asked participants to pick the visualizations that they would be most comfortable sharing. After making their selection, we gave participants an opportunity to orally explain their preferences and to provide any feedback in regards to the visualization designs. Participants' feedback were recorded and later transcribed for analysis.

At the end of the study, participants completed a survey that measures privacy concerns and use of social network sites. Participants were then compensated with a $30 gift card.

7.7   Key Findings

Table 10 shows which location visualization participants picked when asked to choose the visualization that they were most comfortable sharing for each relationship type. Based on these results, we can make three important observations.

	Preferred Location Visualization
Family	map-based, halo marker, general geographic labels (50.0%) map-based, pointing marker, specific geographic labels (33.3%) map-based, pointing marker, general geographic labels (16.7%)
Close Friends	map-based, pointing marker, general semantic labels (83.3%) map-based, halo marker, general semantic labels (16.7%)
Acquaintances	map-based, pointing marker, general semantic labels (67.7%) map-based, halo marker, general geographic labels (33.3%)
Boss/Professors	text-based, general geographic labels (83.3%) map-based, general geographic labels (16.7%)

Table 10. The preferred location visualization for each relationship type, along with the percentage of participants choosing that particular visualization combination as their preferred visualization that they would share with that person.

First, participants unanimously disliked the time-based visualizations for sharing their location history with others. Based on interview feedback, participants reported that sharing their past activities seemed much more privacy-sensitive than sharing only their location information. In the time-based visualization, participants reported that highlighting the temporal aspects of one's location history is more likely to lead others to draw (potentially incorrect) inferences about what they may have been doing at that location.

> P2: *"If they see that I was at some place for a long time, for more than a day, they're going to want to click on it. You know, so that they can find out more. I mean, I would do that if I was looking at someone else's timeline. But actually, you know, I don't really mind people knowing I was home then. It's just that I don't really want to advertise that I was at home for so long?"*

P5: *"You just don't know what kinds of conclusions people are going to jump to when they see how long you spend at certain places. I mean, what if I was at home for a whole week. Maybe I was feeling sick or something. Who knows. But someone could start thinking, gosh he's such a lazy bum. And I don't know if I really want people thinking that...I mean, I can't really control what the timeline will look like to other people, so yeah I just don't feel comfortable sharing it."*

Second, we observed that, in general, participants preferred using general abstractions for their location labels. We saw participants picking this type of label, even in combination with the precise marker style. Upon reflecting on our participants' responses as to why they prefer this type of label, it seems that there are two major privacy concerns influencing participants' decisions about what location information to share. Specifically, participants are concerned about their *physical privacy* and their *image maintenance*.

Participants understood that, by sharing their location information, it might make it easier for others to unexpectedly drop in on them. To prevent these intrusions, participants commented that sharing generic geographic abstractions would make it much harder for others to physically locate them. In other words, participants are worried about their *physical privacy* and are using general abstractions (mostly geographic ones) to provide additional "protection" from being found.

P3: *"I kind of like being able to go someplace and knowing that other people won't be able to find you. If you use this type of visualization [with specific geographic abstractions], you can't really do that anymore. Well, I mean, it kind of makes it super easy for people to bother you whenever they want. So yeah, I'd rather share the other [general geographic abstraction]. If someone really knows you, then yeah they might know where to find you. But those people are OK...like they know you well enough to figure it out, so it probably means they're a really good friend.*

But that's why you need something [general] – so that you don't have to worry about those other people."

P9: *"I guess I can see how spontaneous meetings might be cool. But I think I'd only like it every once and while. I'd rather someone just call me if they want to find me. Giving them this [general] description means they'll still need to ask me to find me. Sure, it's more work for me, but I prefer to know that someone is looking for me?"*

Participants' concerns over *image maintenance* also revealed several interesting privacy concerns relating to information visualizations. For some participants, their location visualizations revealed very clear routine patterns, often associated with shuttling between school/work and home and nothing else. In these cases, location visualization that used the halo marker only served to further emphasize the "routine-ness" of their location history. However, due to the inherent imprecision associated with location sensing, in several cases the participant appears to "move around" when viewing their location information with visualizations that user the pointing marker (even though in reality, their true location never changes). Participants commented that these "micro-motions" provided enough ambiguity that, if one of their friends asked, then they could easily spin a more interesting story.

P10: *"I know these markers kind of reveal where I've been, but I like that the points are kind of scattered around. Sure they're all kind of around this one place, but because there's lots of cooler stuff going around there, I could easily just tell people I was there and not have to seem like that guy who always works all the time. That other one [with the halos] just makes my life seem kind of boring. Which yeah I already know that, but I don't need my friends also thinking that too…haha…"*

Finally, we also observed that, in cases where participants truly wanted to minimize their location sharing (without the option to completely opt out), they nearly always chose the text-based visualization (and always chose to disclose the least descriptive label, i.e., the general geographic abstraction). When asked if text-based visualizations would still be preferable if they were forced to share specific geographic labels (i.e., street address or intersections), participants almost unanimously still chose the text-base visualization (11/12). The reasoning provided by participants is that though the precise makers (on a map-based visualizations) are just as revealing as the street-level descriptions (in the text-based visualizations), the textual natural of the information doesn't solicit further probing of information.

> P4: *"I don't my boss to know anything about where I am. But if I had to share street-level descriptions, then I guess I'd still choose the text[-based visualization]. It just seems like it's more innocent. You can't really immediately see if there's anything sketchy going on unless you really look hard. And bosses are usually busy doing their own thing that they're not going bother with all that extra work."*

> P8: *"Well, my thinking is that when I look at a map and see [a marker], I first think 'do I know that place?' If I don't, then I start thinking about if I know anything near that place. So, even if I don't know that exact place, I still might know about that general area. It's kind of like how I might not know all the bars on some street in South Side, but I know that there are bars there…"*

The last quote from the P8 highlights the extra information that is implicitly embedded in all map-based visualizations. In text-based visualizations, one can look at a particular street address and if there is no immediate knowledge about that address, participants feel fairly confident that most users won't pursue the issue further. However, with maps-based visualizations, people can leverage their general knowledge about the area and starting making inferences about specific

places that they may not know about. It is this potential for drawing inaccurate conclusions that concerns our participants. We refer to this type of privacy concern as relating to *image maintenance.*

7.8   Summary

In this chapter, we presented results from our fourth user study that was designed to exploration how users' perceptions of different location visualizations for sharing one's location history. While we do not claim that we have exhaustively explored the entire design space for visualizing location history, we believe that this study represents a step forward in the understanding of how visual presentation of the same location information can adversely affect end-users' perceptions of privacy and social utility.

By controlling for the same content between visualizations, we found that certain visual elements were more privacy invasive than others. For example, participants felt particular sensitive about sharing how long they spent at a place. In fact, when sharing location history, participants generally felt more sensitive about sharing temporal factors, as opposed to spatial factors. Specifically, out of the four dimensions of location history (spatial coordinates, frequency, duration, and arrival information), participants ranked that they were least comfortable sharing the duration ($\mu=1.6$, $\sigma=0.32$) and arrival information ($\mu=2.9$, $\sigma=0.58$). These self-reported comfort ratings were given using a 5-point Likert scale, where 1=not comfortable and 5=comfortable. This preference is a strong reason for why participants unanimously disliked the time-based visualization for sharing with others. However, many participants did mention that they would appreciate this type of visualization for personal reflection purposes.

When comparing the text- and map-based visualizations, there was a strong preference for sharing a map-based visualization with the different relationship types. In order to provide varying degrees of privacy, the participants opted to use different location labels in their visualizations. This is where it becomes important to consider how to incorporate

location abstractions into visualizations. In particular, we found that users were more comfortable sharing labels that reference general versions of either geographic or semantic location abstractions. The reason for this is that they provided more plausible deniability and, depending on the relationship type, if more deniability is needed, then the more vague abstraction (general geographic abstraction) would be preferred (over the general semantic abstraction).

We also found that users were considering the social utility of the information that they wanted to share. Specifically, we found that participants were giving significant consideration to *impression formation*. In other words, participants were concerned how others would interpret their location visualizations. Thus, we found several participants citing that they wanted to certain visual elements that would enhance their "coolness" or sociability to others. In example of this, would be to pick a visualization that included more visual markers (i.e., the pinpoint markers, as opposed to the halo markers). Participants felt that this would more likely lead to others thinking that they visited more places, as it *appeared* that way due solely to the *quantity* of markers. This type of feedback strongly suggests that there are indeed ways for certain visual manipulations to affect end-user perceptions about issues relating to privacy and social utility for location sharing.

8. Sharing Outcomes: Privacy & Social Utility

In this chapter, we explore our last research question, which is to examine two types of outcomes that can be expected when users engage in social location sharing. In particular, we are interested in seeing what differences are afforded when users choose to share location abstractions, as opposed to more precise descriptions of their location information (see Figure 24). The previous studies thus far have mostly looked at *perceived* concerns about location sharing. In this chapter, we are more interested in *actual* outcomes of location sharing. In particular, we want to understand: 1) whether sharing location abstractions are actually privacy-preserving, and 2) what are the types of social interactions that one could expect from sharing location abstractions.
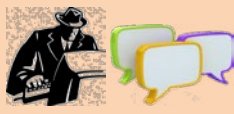
reasoning	configuration	presentation	outcomes
how do users think about location sharing?	how do users specify their sharing preferences?	how do visualizations influence preferences?	what are expected outcomes of sharing?
study 1	study 2 & 3	study 4	study 5 & 6

Figure 24. The four research questions covered in this dissertation. This chapter focuses on the last research question, which looks at two potential outcomes of users' decisions for sharing location abstractions: 1) how privacy preserving users' sharing preferences actually are, and 2) how socially engaging user's location sharing behaviors are.

8.1   Actual Privacy of Location Sharing Decisions

Given that participants factor in privacy concerns when sharing location, our last research question looks at how well participants' decisions *actually* preserve privacy. To do this, we looked at how easily locatable our participant's shared locations were for both purpose-driven and social-driven sharing scenarios.

For each place label that participants shared, we considered how easily locatable they would be if a third party had access to certain resources. The most basic resource would be having a map of the area, or having the ability to conduct local map searches using a tool like Google Maps. The second resource we considered was if the third party had local knowledge of the area (e.g., from being a local resident) or if they had access to a search engine. The third type of resource we considered is if the third party had information about the participant and her routines. One can imagine that this information might be obtained from personally knowing the person or, if more malicious, from stalking the person. As a baseline, we assumed the third party knows at least the participant's current city. For physical stalkers or close friends, this information is obvious. For tech-savvy virtual stalkers, one could imagine that this information could be obtained through basic IP-based geo-location.

We defined a place label as having "revealed" a participant's location if the label means the participant is locatable, i.e., a third party can physically find the person. To run this evaluation, we manually ran the participants' labels through Google's map search (for the map-only resource condition). For the web/local knowledge condition, we used our own knowledge of the local university community combined with a Google web search. We did not expose participants' labels to an actual third-party attacker to ensure participants' data confidentiality.

Each of the resource conditions require different types of labels in order to be found. To be found using only the map resource, the participant must have chosen to share an exact

address, or have disclosed a place label in which the first result of a map-based search query (using only the place label) points to the participant's actual location. To be found using a map with local area knowledge, the participant must have shared a label that can only be resolved with some regional information (e.g., that another resident or community member would know) or be resolved by the first result returned in a search query (using only the place label). To be found using knowledge about routines, the participant must have shared a place label that is easily resolvable based on basic routine information that includes knowing the location of their work and home.
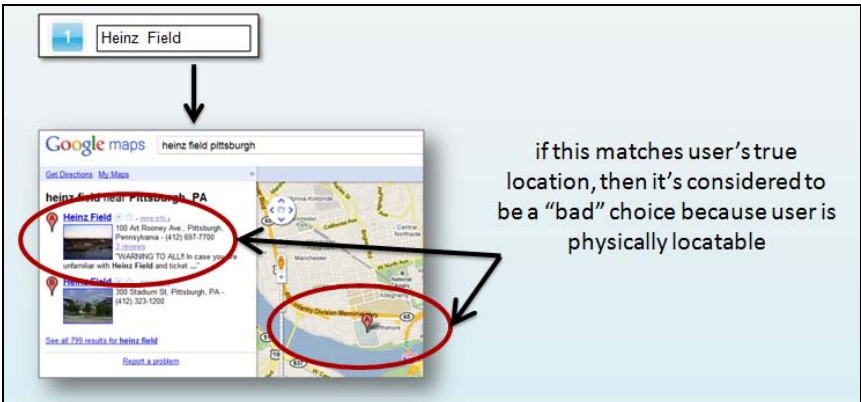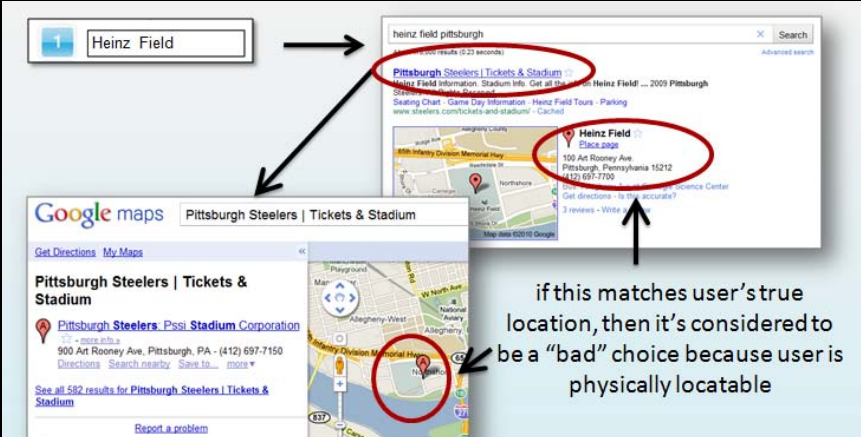
tools	process to "reverse engineer" location labels
google maps	
google search + google maps	
routines + google search + google maps	same as above, but in addition you can: • Leverage personal labels ("home", "work", etc.) • Resolve ambiguous labels ("starbuck") if they are for places that are repeatedly visited

Table 11. A table describing the different ways we leveraged existing knowledge and/or tools to reverse engineer a location label.

Using these definitions, Table 12 shows that, as expected, for purpose-driven location sharing, most of the location disclosures reveal participants' true location. This result is not really disconcerting since participants are aware of who they are sharing their location with. Note that for 9.18% of the place labels which could not be resolved using the three resources, participants were either in-transit or were out of town (and chose to reveal a vague place label).

Available Resources	Purpose-Driven Location Sharing	Social-driven Location Sharing
Map	50.00%	10.20%
Map + Local/Web	62.26%	19.39%
Map + Local/Web + Routines	90.82%	51.02%

Table 12. Percentages of place labels that can lead to physically locating the participant. Organized by resources one might have access to (maps, local info, routines info)

For social-driven location sharing, participants' locations are revealed for at most 51.02% of the labels, when using all three resources. While this percentage is significantly less than purpose-driven sharing (p<0.0001), participants are still locatable for over half of their disclosures. In social network sites, users often unintentionally leak information [Gross and Acquisti, 2005]. In future work, it would be worthwhile to examine if users are aware that the locations they choose to share reveal their true location. Findings from our interviews indicate that sometimes participants reveal their location for impression management or to attract attention. However, since there are also privacy issues to consider, it will be interesting to see whether privacy concerns about the aggregate revelation of location information will lead to changes in users' location sharing decisions over time.

It should be noted that we have adopted a fairly liberal metric for measuring privacy preservation. In particular, we consider someone as "locatable" if they can be found at the building level. However, finding someone at the university student center is not the same level of precision as finding him at a local coffee shop, even though both are considered

building-level granularity. Despite this difference, we believe our analysis provides initial evidence that privacy leaks in social-driven location sharing is an important factor to consider when designing LSAs and is worth further looking into. It is also worth mentioning that many social network sites allow sharing of photos and videos, which can also leak location information. For example, a photo can reveal a well-known landmark, revealing a user's recent whereabouts. This type of information could easily serve as an additional source for locating users. Our initial analysis here shows how different resources can be combined to reveal location information leaks than users may not be aware of.

We also observed that, for social-driven scenarios where they were physically at home, all participants opted to describe their location as "at home", "my home", or "at my apartment." These descriptions were not used for any other locations. Because participants are somewhat predictable in terms of how and when they describe their home, it is important for future social-driven LSAs to have usable privacy controls to limit publicly sharing this data. Otherwise, sites like Please Rob Me [2010] can misuse the data, leaving users open to attacks from malicious users.

8.2   Social Interaction Outcomes of Location Sharing Decisions

In previous chapters, we have explained that, based on past literature, there is sufficient evidence to link context information sharing (such as location sharing) to social awareness to enhanced social bonding. However, we are interested in determining whether there is empirical evidence that will support this claim. We are also interested in determining whether sharing location abstraction in particular leads to different types of social interaction.

To study this issue, we developed a Facebook application that would collect a participant's past status messages, along with their corresponding comment activity. Our intention is to use the comment activity as an indirect measure of the amount and type of

(online) social interaction that occurs between users. While Facebook does support multiple types of communication tools (aside from comment activity), we opted to limit our data collection to comment activity since it is an easily accessible through the Facebook API. Chat messages are not possible to access through the API and inbox messages are much more privacy-intrusive, so we opted to exclude those two communication activities.

We deployed our Facebook application with six undergraduate users and collected 3 months worth of their status messages, spanning from June to August 2010. Participants had, on average, 223.5 friends in their online social network. In total, we collected 3,545 status messages from our participants' wall and 892 comments. We manually sorted through each of these status messages and extracted those that referenced location information and classified them using a similar taxonomy in Chapter 4, where we tag the messages as containing "specific geographic", "general geographic", "semantic", or "hybrid" location descriptions. Of the 3,545 status messages, 12.3% contained location information (436 messages). Very few of these messages contained specific geographic information and those that did were most often a result of third-party applications (e.g., Foursquare) forwarding their information through Facebook.

To measure the amount of social activity, we compared comment activity along three dimensions: the type of users who leave comments (i.e., their relationship to the person who left the comment), the number of comments, and the length of comments. We found that status messages with semantic references were more likely to have marginally more comments ($p < 0.09$) than messages with geographic references. We found no significant differences in terms of the length of comments. However, we did find a significant difference in terms of the types of users who left comments. Status messages that included semantic location descriptions (as opposed to geographic references) were more likely to have comments left by users who have weak social ties to the discloser.

8.3 Summary

In this chapter, we presented results from final two studies that focused on examining the *actual* outcomes of sharing location abstractions. In our previous studies we have focused on more subjective metrics of location sharing and, while it is important to consider these issues when designing LSAs, it is equally important to evaluate them objectively as well.

To do this, we revisited data from our first study to examine how "good" users' decisions are about sharing location abstractions. We specifically wanted to determine how many of the location descriptions could be reverse engineered using simple, publically available tools. We found that, with enough information, a little over half (51%) of the location abstractions could result in the participant being physically locatable. This relatively high percentage indicates that using location abstractions is not a fail-proof way to protect one's privacy. If someone has enough information about the user (e.g., by knowing the user well enough to know about her routines), then it is very likely that he can locate her based on their location sharing behaviors.

However, this dissertation has repeatedly emphasized the importance of considering both the privacy concerns *and* the social utility behind location sharing. Thus, we wanted to quantitatively examine whether it is socially beneficial to share location abstractions. To do this, we examined past Facebook status messages from six participants over a three-month period. We then classified the messages according to whether they contained geographic or semantic location abstractions (or both). We then analyzed these to determine whether there was a correlation between the type of location abstraction used in the status message and the status' comment activity. We found that, while there was no difference in the length of the comments, there was a difference in *who* left comments. In particular, status messages that contained only semantic information, there tended to be more comments and the commenters tended to have weak social ties with the users (i.e., people like acquaintances and casual friends).

By conducting these two types of data analyses, we now have a much better idea of the types of outcomes we could expect if we were to design a LSA that supported location abstractions. We also have quantitative evidence of the types of privacy protection that users need to be aware of, as well as the types of social benefits that one could expect from engaging in social location sharing.

9. Discussion and Design Implications

This dissertation has taken a multi-perspective exploration of disclosure abstractions in social location sharing applications. We have examined multiple types of location sharing applications, including those that simply share one's current location history, as well as those that share one's location history. We have also examined applications that share location information synchronously (similar to services like Google Latitude [2009]), as well as those that share information asynchronously (like the feed-based disclosure model used in Facebook [2004a]). The range of location sharing applications that we have
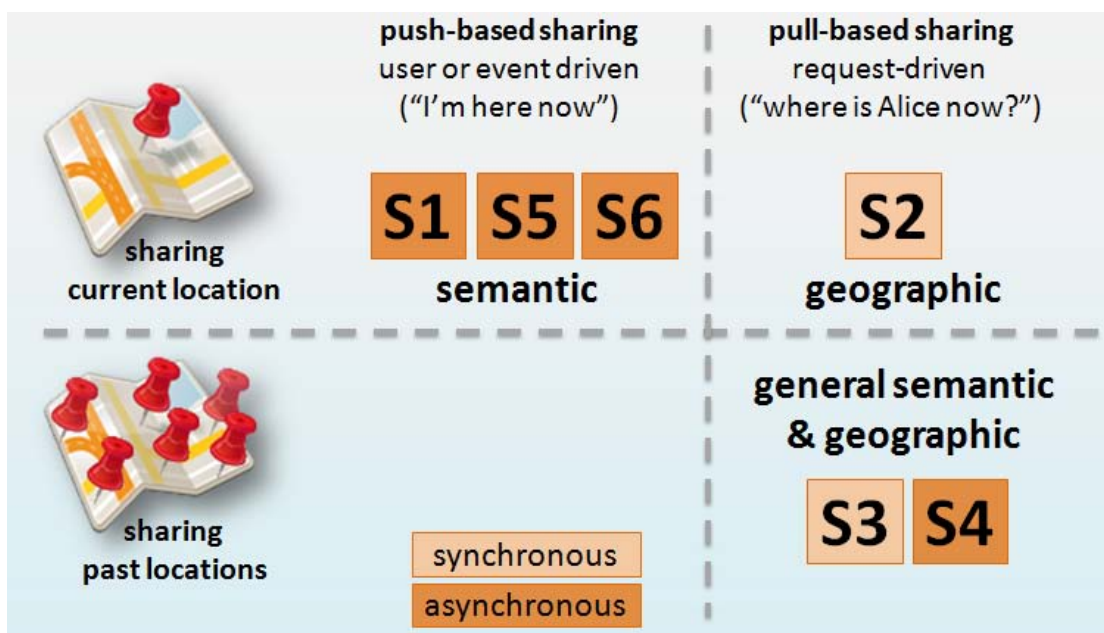

Figure 25. An overview of our six user studies and how they differed in terms of users' preference for location abstractions. These differences could be attributed to the different styles of location sharing (asynchronous vs. synchronous sharing and sharing current vs. past locations).

covered in our studies is shown in Figure 25, along with the preferred location abstractions that we learned from our user studies.

The important take-away is that there is no single answer for which type of location abstraction works best for all location sharing applications. Instead, our results suggest that the abstraction that best balances users' social utility and privacy concerns is dependent on the type of location sharing supported by an application. Applications that focus on sharing only current location information will have different design guidelines than those that focus on sharing past location information.

It is also important to note that, while location applications can technically support a wide range of abstractions (including those that provide a very precise location description), it is often the case that just because an application *can* record some aspect of one's location information, it does not mean that it *should* do so. This becomes an important factor to consider when designing visualizations of users' past location histories. For these types of applications, there are many variables of one's past that could be aggregated and shared, but the service provider must also be sensitive to users' privacy concerns and create visualizations that explicitly *avoid* capturing highly sensitive information. While this may seem unfortunate, in terms of "free" data being "thrown away", it will potentially have long term benefits for the application's membership base, as users will more likely be comfortable sharing their location information with others and less likely to dismiss the application based purely on privacy concerns.

9.1 Specific Design Suggestions for Future LSAs

Based on the preferences outlined in Figure 25, we can make concrete design suggestions for different types of location sharing applications. For example, consider Foursquare [2009]. This type of LSA is more closely aligned to asynchronous sharing of location history, as it supports a feed-based model for disclosing one's past locations. Traditionally, Foursquare requires that users share a specific semantic description (i.e.,

the business name or personal label of a place) and a specific geographic description (i.e., the address of the place). However, according to our results, a more comfortable disclosure model would support *general* semantic and geographic description (see Figure 26). In this case, users would be able to share the genre of the place that they are visiting (e.g., sharing "restaurant" instead of "Imbrie Hall") and the city/neighborhood of the place (e.g., sharing "Hillsboro, OR" instead of the exact street address).

Figure 26. Examples of how we could modify the Foursquare application to support location abstractions.

However, our studies have also alerted to us that it is important to consider the implications of sharing these types of location abstractions. In particular, we know that users could be easily located when given enough location information (at least two weeks worth of data) and enough insider information (e.g., information about one's routines and favorite places). While we did not do any evaluation of whether users are aware of their location information potentially being leaked out, it is most likely that they are not aware that their location sharing behaviors could lead to this outcome.

Because of this, we recommend that LSAs should take active steps to stay aware of these potential inferences, on behalf of the user. The types of processes that we used to reverse engineer the location labels could easily be automated (at least the parts where we relied

on tools like Google Search and Google Maps). If such checks could be incorporated in the location sharing features in LSAs, then the application could proactively suggest to users to describe their location more generally, if they are more concerned about being physically located at a particular place.

On the other hand, we also know that location sharing can lead to more opportunities for social interaction, particularly from weak social ties. LSAs can leverage this information by, again, proactively suggesting to users how they should consider describing their location information. This suggestion would, of course, only apply to users who have already made the decision to share their location information. But for these users, LSAs could use the results of our studies to make sure that the location information that is being shared is done in a way that optimizes the social utility of engaging in such behaviors. For example, LSAs could suggest to a user that, if they want to engage certain types of users (e.g., their casual friends, as opposed to their family members), then it might be more beneficial to share their location information using general geographic descriptions, as opposed to specific geographic descriptions.

LSAs can also increase the amount of location sharing by incorporate more ways to support impression formation. We saw in our study about location visualizations that this concept plays a significant role in users feeling comfortable about sharing their location information. By giving users more ways to manipulate this features (e.g., in our study, we manipulated the marker style and the quantity of markers), LSAs can indirectly influence how useful users perceive location sharing to be. By manipulating the social utility, it is then possible for LSAs to overcome privacy barriers that might have otherwise prevents users from engaging in location sharing.

However, we would be remiss to suggest that future LSAs solely consider social utility in their designs. In all of our studies, participants made it clear that they have significant

privacy concerns about location sharing, particularly when it involves sharing location history. While there are certainly ways to unknowingly convince users to share more information than they might otherwise want to (as seen in our visualization study), it should be the goal of LSA developers to ensure that location sharing is done in a privacy-preserving manner that is usable for end-users. One way to do this is to use location abstractions, which we have shown to be useful in privacy configurations (by simplifying users' privacy policies).

We hope that future LSA developers consider our design suggestions as a way to bridge both end-user privacy concern, as well as social utility issues. Our results strongly suggest that there are ways that can accommodate both sides, without having to completely sacrifice one for the other.

10. Limitations and Future Work

There are a number of topics that fall outside the scope of this dissertation, but could easily be considered as future work that fall within the theme of evaluating how disclosure abstractions can influence end-user privacy concerns.

10.1 Exploring other types of location abstractions

In this dissertation, we have focused on two specific kinds of location abstractions: geographic and semantic. However, there are other types of abstractions that are also worthwhile to consider, depending on the different ways of classifying location information. Further work would need to be done to: 1) examine how to combine different types of abstractions, and 2) determine which abstractions are the most usable for end-users. More abstractions types would certainly increase the complexity and would directly affect how users make decisions about location sharing since there would be a more disclosure options that users could choose from. More studies would need to be done to determine whether the gain in flexibility is worth the cost in complexity.

10.2 Applying the idea of abstractions to other context sources

This dissertation only examined disclosure abstractions for location sharing. However, sharing other types of contextual information can also help to provide and enhance one's social awareness of others. For this work, location sharing seemed like a pragmatic choice given that the current state of mobile technology. In addition, location sharing has

the added benefit of being hierarchical in nature and lends itself well to the concept of location abstractions (i.e., there is a clear sense of when some location label is more or less descriptive than another). For other types of contextual information sharing, particularly those that do not have clear hierarchical qualities, it may be more difficult to incorporate the idea of abstractions. A deeper analysis would be needed to see how our results would translate laterally to others type of data types.

10.3 Running longitudinal studies with location abstractions

Analyzing privacy results from a month-long field trial (e.g., the IMBuddy study) is already a significant improvement over most prior work, which have tended to describe user studies that range from a few days to a week. However, because perceived privacy is an evolving concept, it is important to also consider how people's sharing preferences and behaviors change over a much longer deployment, particularly as users become more exposed to the practice (and hopefully the social benefits) of location sharing.

We also need more empirical evaluation of the design suggestions that we have proposed in Chapter 9. Our dissertation work has laid down the ground work for empirically determining how and why location abstractions should be included in LSAs. As future work, it would be ideal to verify our analyses of our findings in a real-world deployment of a LSA that incorporate our results and design ideas.

Bibliography

148apps. (2010). App Store Metrics. Retrieved October 8, 2010, from
http://148apps.biz/app-store-metrics/?mpage=appcount

ABI Research. (2008). Personal Navigation to Remain Most Popular Lbs Application
over Next Five Years, but Enterprise Applications Will Generate Most Revenue.
Retrieved September 16, 2010, from http://www.abiresearch.com/press/1185

Ackerman, M.S., Cranor, L.F., and Reagle, J. (1999). Privacy in E-Commerce:
Examining User Scenarios and Privacy Preferences. *1st ACM Conference on
Electronic Commerce (E-Commerce '99)*, ACM Press, 1-8.

Acquisti, A. (2004). Privacy in Electronic Commerce and the Economics of Immediate
Gratification. *5th ACM Conference on Electronic Commerce (EC '04)*, New York,
NY, ACM Press, 21-29.

Acquisti, A. (2005). Privacy and Rationality in Individual Decision Making. *IEEE
Security and Privacy,* 3(1), 26-33.

Acquisti, A. (2009). Nudging Privacy: The Behavioral Economics of Personal
Information. *IEEE Security & Privacy,* 7(6), 82-85.

Adams, A. (2000). Multimedia Information Changes the Whole Privacy Ballgame.
*Computers, Freedom &
Privacy (CFP '00)*, 25-32.

Anthony, D., Kotz, D., and Henderson, T. (2007). Privacy in Location-Aware Computing
Environments. *IEEE Pervasive Computing,* 6(4), 64-72.

Aoki, K., and Downes, E.J. (2003). An Analysis of Young People's Use of and Attitudes
toward Cell Phones. *Telematics and Informatics,* 20(4), 349-364.

Aoki, P.M., and Woodruff, A. (2005). Making Space for Stories: Ambiguity in the Design of Personal Communication Systems. *SIGCHI Conference on Human Factors in Computing Systems (CHI 2005) (CHI '05)*, ACM, 181-190.

Apple. (2008). Iphone. Retrieved September 16, 2010, from http://www.apple.com/iphone/apps-for-iphone/

AT&T. (2009). Familymap. Retrieved September 16, 2010, from http://familymap.wireless.att.com/

Avrahami, D., Gergle, D., Hudson, S.E., and Kiesler, S. (2007). Improving the Match between Callers and Receivers: A Study on the Effect of Contextual Information on Cell Phone Interruptions. *Behaviour & Information Technology,* 26(3), 247-259.

Barkhuus, L. (2004). Privacy in Location-Based Services, Concern Vs. Coolness. *Workshop on Location System Privacy at Mobile HCI 2004 (Workshop on Location Privacy, Mobile HCI '04).*

Barkhuus, L., Brown, B., Bell, M., Sherwood, S., Hall, M., and Chalmers, M. (2008). From Awareness to Repartee: Sharing Location within Social Groups. *Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, ACM, 497-506.

Barkhuus, L., and Dey, A.K. (2003). Location-Based Services for Mobile Telephony: A Study of Users' Privacy Concerns. *Interact 2003 (Interact '03)*, 709-712.

Bauer, L., Cranor, L.F., Reeder, R., Reiter, M.K., and Vaniea, K. (2008). A User Study of Policy Creation in a Flexible Access-Control System. *The Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, Florence, Italy, ACM Press, 543-552.

Bellotti, V., and Sellen, A. (1993). Design for Privacy in Ubiquitous Computing Environments. *Third Conference on European Conference on Computer-Supported Cooperative Work (ECSCW '93)*, Kluwer Academic Publishers, 77-92.

Benisch, M., Kelley, P.G., Sadeh, N., Sandholm, T., Cranor, L.F., Drielsma, P.H., et al. (2008). The Impact of Expressiveness on the Effectiveness of Privacy Mechanisms for Location Sharing. Carnegie Mellon University, Institute of Software Research (CMU-ISR-08-141R).

Bentley, F.R., and Metcalf, C.J. (2007). Sharing Motion Information with Close Family and Friends. *SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, ACM, 1361-1370.

Beresford, A.R., and Stajano, F. (2003). Location Privacy in Pervasive Computing. *IEEE Pervasive Computing,* 2(1), 46-55.

Bilton, N. (2010). Facebook Will Allow Users to Share Location. Retrieved September 16, 2010, from http://bits.blogs.nytimes.com/2010/03/09/facebook-will-allow-users-to-share-location/

boyd, d. (2004). Friendster and Publicly Articulated Social Networking. *CHI 2004 (CHI '04)*, ACM, 1279 - 1282.

BrightKite. (2007). Retrieved September 16, 2010, from http://www.brightkite.com

Brown, B., Taylor, A.S., Izadi, S., Sellen, A., Kaye, J., and Eardle, R. (2007). Locating Family Values: A Field Trial of the Whereabouts Clock *Ubicomp 2007 (Ubicomp '07)*, Springer-Verlag, 354-371.

Burrell, J., and Gay, G.K. (2002). E-Graffiti: Evaluating Real-World Use of a Context-Aware System. *Interacting with Computers,* 14(4), 301-312.

Card, S.K., Mackinlay, J., and Shneiderman, B. (1999). Readings in Information Visualization: Using Vision to Think (1st edition ed.). San Francisco: Morgan Kaufmann.

Chen, M., Sohn, T., Chmelev, D., Haehnel, D., Hightower, J., Hughes, J., et al. (2006). Practical Metropolitan-Scale Positioning for Gsm Phones. *Ubicomp 2006 (Ubicomp '06).*

Cheng, Y.-C., Chawathe, Y., LaMarca, A., and Krumm, J. (2005). Accuracy Characterization for Metropolitan-Scale Wi-Fi Localization. *Mobisys 2005 (Mobisys '05).*

Colbert, M. (2001). A Diary Study of Rendezvousing: Implications for Position-Aware Computing and Communications for the General Public. *2001 International ACM SIGGROUP Conference on Supporting Group Work (GROUP '01)*, Boulder, Colorado, ACM Press, 15-23.

Consolvo, S., Smith, I., Matthews, T., LaMarca, A., Tabert, J., and Powledge, P. (2005). Location Disclosure to Social Relations: Why, When, & What People Want to Share. *CHI 2005 (CHI '05)*, 82-90.

Cornwell, J., Fette, I., Hsieh, G., Prabaker, M., Rao, J., Tang, K.P., et al. (2007). User-Controllable Security and Privacy for Pervasive Computing. *8th IEEE Workshop on Mobile Computing Systems and Applications (HotMobile 2007) (HotMobile '07)*.

Counts, S., and Smith, M. (2007). Where Were We: Communities for Sharing Space-Time Trails. *15th Annual ACM International Symposium on Advances in Geographic Information Systems (GIS '07)*, ACM, 1-8.

CTIA. (2008). Wireless Industry Indices: Year-End 2008 Semi-Annual Data Survey. CTIA-The Wireless Association (Report, http://files.ctia.org/pdf/CTIA_Survey_Year-End_2008_Graphics.pdf).

DePaulo, B.M., and Kashy, D.A. (1998). Everyday Lies in Close and Casual Relationships. *Journal of Personality and Social Psychology,* 74(1), 63-79.

Dodgeball. (2009). Retrieved July 13, 2009, from http://www.dodgeball.com

Donath, J., and boyd, d. (2004). Public Displays of Connection. *BT Technology Journal,* 22(4), 71-82.

Dopplr. (2009). Retrieved July 13, 2009, from htpp://www.dopplr.com

Ellison, N.B., Steinfield, C., and Lampe, C. (2007). The Benefits of Facebook "Friends:" Social Capital and College Students' Use of Online Social Network Sites. *Journal of Computer Mediated Communication,* 12(4), article 1.

Erikson, T., Smith, D.N., Kellogg, W.A., Laff, M.R., Richards, J.T., and Bradner, E. (1999). Socially Translucent Systems: Social Proxies, Persistent Conversation, and the Design of 'Babble'. *CHI 1999 (CHI '99)*, 72-79.

Espinoza, F., Persson, P., Sandin, A., Nystrom, H., Cacciatore, E., and Bylund, M. (2001). Geonotes: Social and Navigational Aspects of Location-Based Information Systems. *Ubicomp 2001 (Ubicomp '01)*, Springer-Verlag, 2-17.

Facebook. (2004a). Retrieved September 16, 2010, from http://www.facebook.com

Facebook. (2004b). Facebook Statistics. Retrieved September 16, 2010, from http://www.facebook.com/stats

Facebook. (2007). Leading Websites Offer Facebook Beacon for Social Distribution Retrieved September 16, 2010, from http://www.facebook.com/press/releases.php?p=9166

Facebook. (2010). Places. Retrieved September 16, 2010, from http://www.facebook.com/places/

Farnham, S., and Keyani, P. (2006). Swarm: Hyper Awareness, Micro Coordination, and Smart Convergence through Mobile Group Text Messaging. *39th Annual Hawaii international Conference on System Sciences (HICSS '06)*, IEEE Compter Society.

Federal Communications Commission. (2000). Enhanced 911. Retrieved September 16, 2010, from http://www.fcc.gov/911/enhanced/

Foursquare. (2009). Retrieved September 16, 2010, from http://foursquare.com

Foursquare Grader. (2009). Measure Your Foursquare Mojo. Retrieved October 8, 2010, from http://squaregrader.com/badge/summary

Geonames. (2010). Retrieved September 16, 2010, from http://www.geonames.org/export/ws-overview.html

Glympse. (2008). Retrieved September 16, 2010, from http://www.glympse.com

Glympse. (2009). Retrieved July 13, 2009, from http://www.glympse.com

Google. (2005). Google Maps. Retrieved September 16, 2010, from http://maps.google.com

Google. (2009). Google Latitude. Retrieved September 16, 2010, from http://www.google.com/latitude

Gowalla. (2009). Retrieved September 16, 2010, from http://gowalla.com

Grinter, R.E., and Eldridge, M. (2001). Y Do Tngrs Luv 2 Txt Msg? *Seventh European Conference on Computer-*
Supported Cooperative Work (ECSCW '01), 219-238.

Groovr. (2009). Retrieved July 13, 2009, from http://www.groovr.com

Gross, R., and Acquisti, A. (2005). Information Revelation and Privacy in Online Social Networks. *Workshop on Privacy in the Electronic Society (WPES 2005) (WPES '05)*.

Grudin, J., and Horvitz, E. (2003). Presenting Choices in Context:Approaches to Information Sharing. *Workshop on Ubicomp communities: Privacy as Boundary Negotiation, Ubicomp 2003 (Workshop on Ubicomp Privacy, Ubicomp '03).*

Gruteser, M., and Grunwald, D. (2003). Anonymous Usage of Location-Based Services through Spatial and Temporal Cloaking. *1st international conference on Mobile systems, applications and services (Mobisys 2003) (Mobisys '03)*, ACM, 31-42.

Gruteser, M., and Liu, X. (2004). Protecting Privacy in Continuous Location-Tracking Applications. *IEEE Security and Privacy,* 2(2), 28-34.

Harper, R. (1995). Why People Do and Don't Wear Active Badges: A Case Study *Computer Supported Cooperative Work (CSCW) (CSCW '95)*, Springer-Verlag, 297-318.

Harrison, S., and Dourish, P. (1996). Re-Place-Ing Space: The Roles of Place and Space in Collaborative Systems. *ACM Conference on Computer Supported Cooperative Work (CSCW '96)*, Boston, Massachusetts, ACM Press, 67-76.

Healey, C.G., K.S., B., and Enns, J.T. (1995). Visualizing Real-Time Multivariate Data Using Preattentive Processing *ACM Transactions on Modeling and Computer Simulation,* 5(3), 190-221.

Hindus, D., Mainwaring, S.D., Leduc, N., Hagström, A.E., and Bayley, O. (2001). Casablanca: Designing Social Communication Devices for the Home. *SIGCHI Conference on Human Factors in Computing Systems (CHI '01)*, ACM, 325-332.

Hindus, D., and Schmandt, C. (1992). Ubiquitous Audio: Capturing Spontaneous Collaboration. *1992 ACM Conference on Computer-Supported Cooperative Work (CSCW '92)*, ACM, 210-217.

Hong, J.I. (2003). Privacy and Security in the Location-Enhanced World Wide Web. *Workshop on Privacy at Ubicomp 2003 (Workshop on Privacy at Ubicomp '03).*

Hong, J.I. (2004). Privacy Risk Models for Designing Privacy-Sensitive Ubiquitous Computing Systems. *5th conference on Designing interactive systems: processes, practices, methods, and techniques (DIS '04)*, ACM, 91-100.

Hong, J.I. (2005a). An Architecture for Privacy-Sensitive Ubiquitous Computing. University of California at Berkeley, Berkeley.

Hong, J.I. (2005b). An Architecture for Privacy-Sensitive Ubiquitous Computing. Unpublished Ph.D. Thesis, University of California at Berkeley, Berkeley.

Hsieh, G., Tang, K.P., Low, W.Y., and Hong, J.I. (2007). Field Deployment of Imbuddy: A Study of Privacy Control and Feedback Mechanisms for Contextual Im. *Ubicomp 2007 (Ubicomp '07)*, Springer-Verlag, 91-108.

Hull, R., Kumar, B., Lieuwen, D., Patel-Schneider, P.F., Sahuguet, A., Varadarajan, S., et al. (2004). Enabling Context-Aware and Privacy-Conscious User Data Sharing. *2004 IEEE International Conference on Mobile Data Management (MDM'04) (MDM '04)*, IEEE Computer Society, 187-198.

Iachello, G., and Hong, J.I. (2007). End-User Privacy in Human-Computer Interaction. *Foundations and Trends in HCI,* 1(1), 1-137.

Iachello, G., Smith, I., Consolvo, S., Abowd, G.D., Hughes, J., Howard, J., et al. (2005). Control, Deception, and Communication: Evaluating the Deployment of a Location-Enhanced Messaging Service. *Ubicomp 2005 (Ubicomp '05)*, 213-231.

Iachello, G., Smith, I., Consolvo, S., Chen, M., and Abowd, G.D. (2005). Developing Privacy Guidelines for Social Location Disclosure Applications and Services. *2005 Symposium on Usable Privacy and Security (SOUPS) (SOUPS '05)*, ACM, 65-76.

Jiang, X. (2002). Approximate Information Flows: Socially-Based Modeling of Privacy in Ubiquitous Computing *Ubicomp 2002 (Ubicomp '02)*, Springer-Verlag, 176-193.

Jones, R., Kumar, R., Pang, B., and Tomkins, A. (2007). "I Know What You Did Last Summer": Query Logs and User Privacy. *Sixteenth ACM Conference on Conference on Information and Knowledge Management (CIKM '07)*, Lisbon, Portugal, ACM Press, 909-914.

Jung, Y., Persson, P., and Blom, J. (2005). Dede: Design and Evaluation of a Context-Enhanced Mobile Messaging System. *CHI 2005 (CHI '05).*

Kaasinen, E. (2003). User Needs for Location-Aware Mobile Services. *Personal and Ubiquitous Computing,* 7(1), 70-79.

Keyani, P., and Farnham, S. (2005). Swarm: Text Messaging Designed to Enhance Social Coordination In R. Harper, L. Palen & A. Taylor (Eds.), *The inside Text: Social, Cultural and Design Perspectives on Sms* (Vol. 4, pp. 287-304): Springer-Verlag.

Khalil, A., and Connelly, K. (2006). Context-Aware Telephony: Privacy Preferences and Sharing Patterns. *2006 20th Anniversary Conference on Computer Supported Cooperative Wor (CSCW '06)*, ACM, 469-478.

Kumaraguru, P., and Cranor, L.F. (2005). Privacy Indexes: A Surey of Westin's Studies. Carneigie Mellon University, Institute of Software Research (Technical Report: CMU-ISRI-05-138).

LaMarca, A., Chawathe, Y., Consolvo, S., Hightower, J., Smith, I.E., Scott, J., et al. (2005). Place Lab: Device Positioning Using Radio Beacons in the Wild. *Third International Conference on Pervasive Computing (Pervasive 2005) (Pervasive '05)*, 116-133.

Lampe, C., Ellison, N., and Steinfield, C. (2006). A Face(Book) in the Crowd: Social Searching Vs. Social Browsing. *20th Anniversary Conference on Computer Supported Cooperative Work (CSCW '06)*, Banff, Alberta, Canada, ACM Press, 167-170.

Langheinrich, M. (2002). A Privacy Awareness System for Ubiquitous Computing Environments *Ubicomp 2002 (Ubicomp '02)*, Springer-Verlag, 315-320.

Laurier, E. (2001). Why People Say Where They Are During Mobile Phone Calls. *Environment and Planning D: Society and Space,* 19(Pion), 485-504.

Lederer, S., Hong, J.I., Dey, A.K., and Landay, J.A. (2004). Personal Privacy through Understanding and Action: Five Pitfalls for Designers. *Personal and Ubiquitous Computing,* 8(6), 440-454.

Lederer, S., Hong, J.I., Dey, A.K., and Landay, J.A. (2005). Five Pitfalls in the Design for Privacy. In S. Garfinkel & L.F. Cranor (Eds.), *Security and Usability* (pp. 421-445).

Lederer, S., Hong, J.I., Jiang, X., Dey, A.K., Landay, J.A., and Mankoff, J. (2003). Towards Everyday Privacy for Ubiquitous Computing. University of California, Berkeley, EECS Department (Technical Report: UCB/CSD-03-1283).

Lederer, S., Mankoff, J., and Dey, A.K. (2003). Who Wants to Know What When? Privacy Preference Determinants in Ubiquitous Computing. *SIGCHI 2003 Extended Abstracts on Human Factors in Computing Systems (CHI '03)*, ACM, 724-725.

Liedtke, M. (2007). Facebook Backpedals from New Ad System. Retrieved September 16, 2010, from http://www.msnbc.msn.com/id/22117173/

Lin, J., Xiang, G., Hong, J.I., and Sadeh, N. (2010). Modeling People's Place Naming Preferences in Location Sharing. *12th ACM International Conference on Ubiquitous Computing (Ubicomp '10)*, ACM Press, to appear.

Loopt. (2005). Retrieved September 16, 2010, from http://www.loopt.com

Ludford, P.J., Frankowski, D., Reily, K., Wilms, K., and Terveen, L. (2006). Because I Carry My Cell Phone Anyway: Functional Location-Based Reminder Applications. *SIGCHI Conference on Human Factors in Computing Systems (CHI '06)*, ACM, 889-898.

Ludford, P.J., Priedhorsky, R., Reily, K., and Terveen, L. (2007). Capturing, Sharing, and Using Local Place Information. *SIGCHI Conference on Human Factors in Computing Systems (CHI '07)*, 1235-1244.

Lynch, C.G. (2007). Wake-up Call in Facebook-Beacon Controversy. Retrieved September 16, 2010, from http://www.pcworld.com/businesscenter/article/140372/wakeup_call_in_facebook beacon_controversy.html

MacEachren, A. (1995). How Maps Work: Guilford Press.

Marmasse, N., and Schmandt, C. (2000). Location-Aware Information Delivery with Commotion *2nd international Symposium on Handheld and Ubiquitous Computing*, Springer-Verlag, 157-171.

Marmasse, N., Schmandt, C., and Spectre, D. (2004). Watchme: Communication and Awareness between Members of a Closely-Knit Group. *Ubicomp 2004 (Ubicomp '04)*.

May, A., Ross, T., Bayer, S., and Tarkiainen, M. (2003). Pedestrian Navigation Aids: Information Requirements and Design Implications. *Personal Ubiquitous Computing,* 7(6), 331-338.

Meyer, D. (2008). Boom Predicted for Gps-Enabled Handsets. Retrieved September 16, 2010, from http://news.cnet.com/Boom-predicted-for-GPS-enabled-handsets/2100-1039_3-6226211.html

Microsoft. (2000). Mappoint. Retrieved September 16, 2010, from http://www.microsoft.com/mappoint/default.mspx

Microsoft Research. (2009). Slam: Social Location Annotation Mobile. Retrieved July 13, 2009, from http://www.msslam.com/About.aspx

Mokbel, M.F.C., C., and Aref, W.G. (2006). The New Casper: Query Processing for Location Services without Compromising Privacy. *32nd International Conference on Very Large Data Bases*, Seoul, Korea, VLDB Endowment, 763-774.

Moore, R.J. (2010). Foursquare Is Five Times Larger Than Gowalla and Growing 75 Percent Faster Every Day. Retrieved September 16, 2010, from http://techcrunch.com/2010/07/07/foursquare-gowalla-stats/

Nagel, K., Kidd, C.D., O'Connell, T., Dey, A.K., and Abowd, G.D. (2001). The Family Intercom: Developing a Context-Aware Audio Communication System. *Ubicomp 2001 (Ubicomp '01)*, Springer-Verlag, 176-183.

Nardi, B., Whittaker, S., and Bradner, E. (2000). Interaction and Outeraction: Instant Messaging in Action. *ACM Conference on Computer-Supported Cooperative Work (CSCW) (CSCW '00)*, 79-88.

Nguyen, D.H., and Mynatt, E.D. (2002). Privacy Mirrors: Understanding and Shaping Socio-Technical Ubiquitous Computing Systems. Georgia Institute of Technology, GVU (Technical Report: GIT-GVU-02-16).

O'Neill, N. (2010). Twitter Nears Facebook's Daily Status Update Volume. Retrieved September 16, 2010, from http://www.allfacebook.com/2010/02/twitter-facebook-status/

Oster, S. (2005). Java Aimbot. Retrieved September 16, 2010, from http://sourceforge.net/projects/jaimbot

Oulasvirta, A., Petit, R., Raento, M., and Tiitta, S. (2007). Interpreting and Acting on Mobile Awareness Cues. *Human-Computer Interaction,* 22(1), 97-135.

Oulasvirta, A., Raento, M., and Tiitta, S. (2005). Contextcontacts: Re-Designing Smartphone's Contact Book to Support Mobile Awareness and Collaboration. *7th*

International Conference on Human Computer interaction with Mobile Devices & Services (MobileHCI '05), ACM, 167-174.

Parr, B. (2010). The Rise of Foursquare in Numbers. Retrieved October 8, 2010, from http://mashable.com/2010/03/12/foursquare-stats/

Patil, S., and Lai, J. (2005). Who Gets to Know What When: Configuring Privacy Preferences in an Awareness Application. *SIGCHI Conference on Human Factors in Computing Systems (CHI '05)*, Portland, OR, ACM Press, 101-110.

Plazes. (2004). Retrieved September 16, 2010, from http://www.plazes.com

Please Rob Me. (2010). Retrieved September 16, 2010, from http://pleaserobme.com/

Radar. (2009). Retrieved July 13, 2009, from http://radar.net

Reilly, D., Dearman, D., Ha, V., Smith, I., and Inkpen, K. (2006). "Need to Know": Examining Information Need in Location Discourse. *Pervasive 2006 (Pervasive '06)*, 33-49.

Reily, K., Ludford, P.J., and Terveen, L. (2008). Sharescape: An Interface for Place Annotation. *5th Nordic Conference on Human-Computer Interaction (NordCHI '08)*, ACM, 326-333.

Resnick, P. (2001). Beyond Bowling Together: Sociotechnical Capital. In J. Carroll (Ed.), *Hci in the New Millennium*: Addison-Wesley.

RunKeeper. (2008). Retrieved September 16, 2010, from http://www.runkeeper.com/

Sadeh, N., Hong, J.I., Cranor, L., Fette, I., Kelley, P., Prabaker, M., et al. (2009). Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application. *Personal and Ubiquitous Computing,* Forthcoming, To appear.

Sadeh, N., Hong, J.I., Cranor, L., Fette, I., Kelley, P.G., Prabaker, M., et al. (2008). Understanding and Capturing People's Privacy Policies in a Mobile Social Networking Application. *Personal and Ubiquitous Computing,* 13(6), 401-412.

Schilit, B.N., LaMarca, A., Borriello, G., Griswold, W.G., McDonald, D., Lazowska, E., et al. (2003). Challenge: Ubiquitous Location-Aware Computing and The "Place Lab" Initiative. *1st ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH '03)*, San Diego, California, ACM Press, 29-35.

Siegler, M. (2010). Twitter Joins the Place Race - Foursquare, Gowalla Come Along for the Ride. Retrieved September 16, 2010, from http://techcrunch.com/2010/06/14/twitter-foursquare-gowalla/

Skyhook Wireless. (2003). Location Apps. Retrieved September 16, 2010, from http://skyhookwireless.com/locationapps/

Skyhook Wireless. (2009a). Location Aware App Report: Review of Location-Aware Apps from the Apple, Blackberry, Android, Nokia and Palm App Stores (Report, http://www.locationrevolution.com/stats/skyhookjulyreport.pdf).

Skyhook Wireless. (2009b). Location Aware App Report: Review of Location-Aware Apps from the Iphone, Blackberry and Android App Stores (Report, http://www.locationrevolution.com/stats/skyhookaprilreport.pdf).

Smith, I., Consolvo, S., LaMarca, A., Hightower, J., Scott, J., Sohn, T., et al. (2005a). Social Disclosure of Place: From Location Technology to Communication Practices. *Pervasive 2005 (Pervasive '05)*, 134-151.

Smith, I., Consolvo, S., LaMarca, A., Hightower, J., Scott, J., Sohn, T., et al. (2005b). Social Disclosure of Place: From Location Technology to Communication Practices *(Pervasive '05)*, 134-151.

Sohn, T., Li, K.A., Griswold, W.G., and Hollan, J.D. (2008). A Diary Study of Mobile Information Needs. *Twenty-Sixth Annual SIGCHI Conference on Human Factors in Computing Systems (CHI '08)*, Florence, Italy, ACM Press, 433-442.

Tang, J.C., Yankelovich, N., Begole, J., Kleek, M., Li, F., and Bhalodia, J. (2001). Connexus to Awarenex: Extending Awareness to Mobile Users. *CHI 2001 (CHI '01)*, ACM, 221-228.

Tang, K.P., Keyani, P., Fogarty, J., and Hong, J.I. (2006). Putting People in Their Place: An Anonymous and Privacy-Sensitive Approach to Collecting Sensed Data in Location-Based Applications. *SIGCHI Conference on Human Factors in Computing Systems 2006 (CHI '06)*, ACM, 93-102.

Tang, K.P., Lin, J., Hong, J.I., Siewiorek, D.P., and Sadeh, N. (2010). Rethinking Location Sharing: Exploring the Implications of Social-Driven Vs. Purpose-Driven Location Sharing. *12th ACM International Conference on Ubiquitous Computing (Ubicomp '10)*, Copenhagen, Denmark, ACM Press, 85-94.

Treu, G., Fuchs, F., and Dargatz, C. (2007). Implicit Authorization for Accessing Location Data in a Social Context. *Second International Conference on Availability, Reliability and Security (ARES '07)*, IEEE, 263-277.

Tsai, J., Kelley, P., Cranor, L., and Sadeh, N. (2009). Location-Sharing Technologies: Privacy Risks and Controls. *The 37th Research Conference on Communication,Information, and Internet Policy (TPRC '09)*.

Tsai, J., Kelley, P., Drielsma, P., Cranor, L.F., Hong, J.I., and Sadeh, N. (2009). Who's Viewed You?: The Impact of Feedback in a Mobile Location-Sharing Application. *27th international Conference on Human Factors in Computing Systems (CHI '09)*, ACM, 2003-2012.

Twitter. (2006). Retrieved September 16, 2010, from http://twitter.com

Twitter. (2009a). Retrieved July 13, 2009, from http://twitter.com

Twitter. (2009b). Location, Location, Location. Retrieved September 16, 2010, from http://blog.twitter.com/2009/08/location-location-location.html

Twitter. (2010). Annotations Overview. Retrieved September 16, 2010, from http://apiwiki.twitter.com/Annotations-Overview

Van Grove, J. (2010). Foursquare Nearing 1 Million Checkins Per Day. Retrieved September 16, 2010, from http://mashable.com/2010/05/28/foursquare-checkins/

Varshavsky, A., Chen, M., Lara, E.d., Froehlich, J., Haehnel, D., Hightower, J., et al. (2006). Are Gsm Phones the Solution for Localization? *7th IEEE Workshop on Mobile Computing Systems and Applications (HotMobile 2006) (HotMobile '06)*.

Ware, C. (1999). Information Visualization: Perception for Design (2nd Edition ed.). San Francisco: Morgan Kaufman.

Weilenmann, A. (2003). "I Can't Talk Now: I'm in a Fitting Room": Formulating Availability and Location in Mobile Phone Conversations. *Environment and Planning A,* 35(Pion), 1589-1605.

Weiser, M. (1991). The Computer for the 21st Century. *Scientific American,* 265(3), 94-104.

Wellman, B., Haase, A.Q., Witte, J., and Hampton, K. (2001). Does the Internet Increase, Decrease, or Supplement Social Capital? Social Networks, Participation, and Community Commitment. *American Behavioral Scientist,* 45(3), 436-455.

Westin, A. (1991). Harris-Equifax Consumer Privacy Survey. Equifax Inc.

Whalen, J. (1995). You're Not Paranoid: They Really Are Watching You. *Wired Magazine,* 3(3)*,* 85-95.

Where. (2009). Retrieved July 13, 2009, from http://www.where.com

Wikipedia. (2001a). Retrieved September 16, 2010, from http://www.wikipedia.org

Wikipedia. (2001b). Local Search (Internet). Retrieved Nov 22, 2010, from http://en.wikipedia.org/wiki/Local_search_%28Internet%29

Yelp. (2004). Retrieved September 16, 2010, from http://www.yelp.com

Zahradnik, F. (2009). Gps-Enabled Phone Market to Grow 6.4 Percent in 2009 Retrieved September 16, 2010, from http://gps.about.com/b/2009/01/22/gps-enabled-phone-market-to-grow-64-in-2009.htm

Zhou, C., Ludford, P., Frankowski, D., and Tervee, L. (2005). Talking About Place: An Experiment in How People Describe Places (2005). *Third International Conference on Pervasive Computing (Pervasive '05).*

Zuckerberg, M. (2006). An Open Letter from Mark Zuckerberg. Retrieved September 16, 2010, from http://blog.facebook.com/blog.php?post=2208562130

Zuckerberg, M. (2007). Thoughts on Beacon. Retrieved September 16, 2010, from http://blog.facebook.com/blog.php?post=7584397130